Essays in Educational and Intergenerational Inequality

Ursula Mattioli Mello

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in

Economics

Universidad Carlos III de Madrid

Advisors:
Jan Stuhler and Ricardo Mora

July 2019

A Deus, pela vida e por tantas maravilhas.

Às minhas origens, à minha família, em especial aos meus avós, por todo esforço para que que eu chegasse até aqui.

Aos meus pais e à Ingrid, por todo amor que existe nesse mundo.

Ao Tomás, por caminhar ao meu lado e por construir sonhos comigo.

Aos meus futuros filhos, que, mesmo do mundo mental, já movem forças inigualáveis dentro de mim.

Acknowledgements

"... to reach this stage (of great achievements), he must create many inner defenses, become valiant in the face of all adverse circumstances, and come to know that struggle is the law of life, which must be confronted not once but a thousand times, and not with vacillation but by being fully conscious of what is inescapable." From the Logosophical Bibliography

These past six years were, definitely, the most challenging and enriching of my life so far. The journey towards a PhD thesis led me to great achievements, but also presented many struggles. Interestingly, most of these battles happened inside of me. During this time, while learning to be a research economist, I had to face fears, develop inner defenses and work even harder to build key virtues for academic and life success: patience, bravery, humility, self-confidence, kindness towards myself and others. Today, I am very happy to have produced a thesis about topics I care about. Yet, there are no words to describe how joyful I feel to realize how much I have grown and changed. Of course, there is no way I could have gotten this far without the help of many.

I am extremely grateful to my supervisors, Jan Stuhler and Ricardo Mora. Thank you for an impeccable guidance that went far beyond improving my research output. Thank you for believing in me and in my work more than I did myself. Thank you for all the time devoted to me. Thank you for giving me the perfect balance between the freedom I needed to discover my own path inside academia and the incentives that helped me expand my limits and persist. You were true examples of how academic excellence can meet an extremely generous and kind heart. I hope that, one day, I get to be half the mentor you were to me.

I am also grateful to the faculty, my fellow PhD friends and the administrative personnel at UC3M. You have created a warm and helpful environment, in which I could grow and learn daily. I am especially thankful to Matilde Machado and Nazarii Saalish, who generously took their time to review my dissertation chapters and give me great advice on how to improve my work. Moreover, I am grateful to the UC3M Applied Faculty and PhD students, whose comments, generosity and example helped build my skills and motivate my work. Thank you for being such great peers and role models. I would also like to say a special thank you to Pedro, Bea, Ana and Conchi, who kindly took care of me during the first two very challenging years of this experience.

I am grateful to my friends who accompanied me during this, in addition to many other experiences of my life. To my friends who happen to be economists, thank you for all the insightful conversations that have inspired many of my research questions. To my forever-friends, thank you for always receiving me with your whole heart after so many months apart. To my friends from the Logosophical Foundation, thank you for being my family in Spain, for filling my life with joy and for helping me find purpose within myself.

Finally, I dedicate this thesis to my loving family, especially, to my parents, Carlos and Sandra, to my sister, Ingrid, and to my love, Tomás. Tomás, thank you for being my best friend, my favorite economist and my life partner through all these six years. Your loyalty, your grit, your insightfulness and your strength do not seize to surprise and inspire me. Thank you for walking by my side every day. Daddy, mommy e sis, thank you for always loving, supporting and trusting me. Thank you for being so close, no matter how far. Thank you for your example of happiness, courage, determination and goodness of heart. Your love makes me brave and my home will always be wherever you are. *Muito obrigada por tudo. Vocês são o melhor de mim. Eu amo muito vocês e espero, um dia, ser para meus filhos tudo o que vocês são para mim.*

Other Research Merits

This thesis has been possible thanks to the financial support of the Spanish Ministry of Education and Culture, grant number FPU14/07116.

Abstract

In this dissertation, I study different dimensions of socioeconomic inequality. My objective is to produce empirical evidence to help governments design better policies in the fields of education and the labor market. The two first chapters are focused on educational inequality and the effect of policies that aim to tacke this issue, while the third chapter aims to present useful solutions to an important conceptual issue for measuring intergenerational mobility.

In the first chapter, '*Affirmative Action, Centralized Admissions and Inequality in Access to Higher Education: Evidence from Brazil*', I analyze how two major reforms, introduced to democratize access to public higher education in Brazil, impacted enrollments of students from a low-socioeconomic status. The first policy centralized applications in a nationwide platform and the second expanded affirmative action quotas to a uniform share of fifty percent of all vacancies offered by each major and institution. Their progressive adoption generates cross-sectional and time variation, allowing the separate identification of their causal effects. Results show that the affirmative action reform increases enrollments of public school, black and low-income students, while the centralized admission system acts in the opposite direction, decreasing their participation. Moreover, the interaction between both policies has a positive and significant effect on enrollments of the vulnerable groups. I then shed light on some mechanisms behind these results. I find that centralization disproportionately increases enrollments of high-SES out-of-state students in the least prestigious degrees, crowding-out low-SES students with mobility constraints. On the other hand, the expansion of affirmative action does not only mechanically improve equity, but also changes application behavior.

In the second chapter, '*Does Affirmative Action in Undergraduate Education Impact High Schools?*', I delve into the analysis of unintended consequences of affirmative action initiatives in higher education, which have been implemented in different countries to improve access of vulnerable groups and to reduce inequality in educational attainment. A growing empirical literature has investigated how such policies impact college students' outcomes and pre-college human capital accumulation. Yet, little is known about how they affect students' choice of high school and, consequently, school quality and peer interaction. I study this question in the context of Brazil, one of the most unequal countries in the world, and where the government

approved, in 2012, the "Quota Law (*QL*)". It established that fifty percent of all vacancies in each major and federal higher education institution, including some of the best universities in the country, has to be reserved to students that attended secondary education integrally in a public school. I show that the adoption of *QL* increases strategic mobility from private to public schools by 29 percent and that the movers come disproportionately from low-SES and low-quality private schools. Nevertheless, this exogenous influx of private school students increases public school quality, while it also raises inequality within the public school system.

Finally, in the third chapter, '*Correction Methods for Intergenerational Mobility Estimates*', co-authored with Martin Nybom and Jan Stuhler, we study another dimension of socioeconomic inequality: the transmission of economic status between generations. The estimation of standard measures of intergenerational mobility ideally requires the complete income history for two generations to determine their lifetime incomes. However, empirical applications are typically based on snapshots of income over a limited number of observations in the life cycle. If those snapshots do not mimic lifetime outcomes, the estimates are subject to attenuation and lifecycle bias. The literature has followed two different strategies to address this problem. The first models the income processes itself, the second the relation between annual and lifetime incomes over the life cycle. In this paper, we use uniquely long income series from Sweden to study how well these methods approximate the intergenerational elasticity of income. All methods are biased to some degree, because neither accounts for three key components of the income process: (i) income growth explained by observable characteristics, (ii) transitory noise, and (iii) unexplained income growth that correlates within families. We propose a lifecycle estimator that addresses all three components, and which can improve estimates of the intergenerational elasticity in a wide range of settings.

Contents

1 Affirmative Action, Centralized Admissions and Inequality in Access to Higher Education: Evidence from Brazil 1

1.1 Introduction . . . 1
1.2 Motivation . . . 7
1.2.1 Brazilian Higher Education System . . . 7
1.2.2 Centralization of Admissions . . . 8
1.2.3 Affirmative Action . . . 10
1.3 Data . . . 11
1.4 Empirical Methodology . . . 14
1.5 Main Results . . . 15
1.5.1 Enrollments of Low-SES students . . . 15
1.5.2 Heterogeneity of Results . . . 16
1.6 Potential Mechanisms . . . 17
1.6.1 Why centralizing admissions crowds out low-SES students from the public higher education? . . . 18
1.6.2 Is the effect of AA on enrollments purely mechanical or is there evidence of any effect on students' behavior? . . . 20
1.6.3 Why does the interaction between SISU and AA creates an additional enrollment-effect for vulnerable groups? . . . 21
1.7 Robustness of Results and Extensions . . . 22
1.7.1 Spillovers . . . 22
1.7.2 Dynamics . . . 24
1.7.3 Sample Selection . . . 25
1.7.4 Additional Specifications . . . 26
1.8 Concluding Remarks . . . 27
1.9 Figures and Tables . . . 30
1.A Additional Figures and Tables . . . 44

2 Does Affirmative Action in Undergraduate Education Impact High Schools? 53

2.1 Introduction . . . 53
2.2 Institutional Background . . . 57
2.2.1 Brazilian Educational System . . . 57
2.2.2 The Quota Law in Higher Education . . . 58
2.3 Data and Variables . . . 59
2.4 Strategic High School Choice . . . 62
2.4.1 Empirical Strategy . . . 62
2.4.2 Results . . . 63
2.5 Effect on the High School System . . . 65
2.6 Concluding Remarks . . . 67
2.7 Figures and Tables . . . 69

3 Correction Methods for Intergenerational Mobility Estimates 86

3.1 Introduction . . . 86
3.2 Data . . . 91
3.3 The Income Process: Swedish and Simulated Data . . . 92
3.4 Correction Methods in the Intergenerational Literature . . . 96
3.4.1 Modelling Errors-in-Variables . . . 96
3.4.2 Modelling the Income Process . . . 100
3.5 New Correction Methods . . . 108
3.5.1 The Lifecycle Estimator . . . 108
3.5.2 The Standardized Errors-in-Variables Model . . . 114
3.6 Concluding Remarks . . . 117
3.A Additional Figure and Tables . . . 119

List of Tables

1.1 Expansion of SISU in Brazilian Undergraduate Education 32
1.2 Expansion of Affirmative Action in Brazilian Undergraduate Institutions . . 33
1.3 Descriptive Statistics of Incoming Students in Public Higher Education Institutions 34
1.4 Placebo Experiment with Lead Variables 35
1.5 Effect of SISU and AA on Enrollments of Low-SES Students 35
1.6 Effect of AA and SISU on Enrollments by Quartile of Baseline Share of Public-School Enrollment in Baseline 36
1.7 Effect on the Number of Applications and Enrollments of Out-of-State Students 37
1.8 Effect of AA and SISU on Average Grades of Incoming Students 37
1.9 Heterogeneity of Effect by Quartile of Competitiveness of Degree 38
1.10 Heterogeneity of Effect on Enrollment of Out-of-State Students by Quartile of Competitiveness of Degree and by Type of High School 39
1.11 Heterogeneity of Effect on Grades by Quartile of Competitiveness of Degree . 40
1.12 Spillover of Treatment Effects on Public Higher Education Market 41
1.13 Yearly Effects 42
1.14 Treatment status and Sample selection 43
B.1 Treatment and Control Units 45
B.2 Robustness of Placebo Experiment 46
B.3 Robustness of Spillover Measures 47
B.4 Robustness - Different Samples of Institutions 48
B.5 Robustness - Different Sample Selections 49
B.6 Robustness - Results Collapsed at Program Level 50
B.7 Robustness - Collapsed at Program Level and Different Local Controls . . . 51
B.8 Robustness - SISU Treatment Collapsed at Institutional Level 52

2.1 Descriptive Statistics 76
2.2 Regression on Treatment Status in 2011 77

2.3 Moves from Private to Public School . . . 78
2.4 Moves from Public to Private School . . . 79
2.5 Moves from Private to Public School - Non-white Students . . . 80
2.6 Moves from Private to Public School - White and Asian Students . . . 81
2.7 Moves from Private to Public School - By School Socioeconomic Level . . . 82
2.8 Moves from Private to Public School - By School Quality . . . 83
2.9 Distribution of $SharePriv_{pm,2010}$ . . . 84
2.10 Effect on $SharePriv_{pmy}$ . . . 84
2.11 Effect on School Quality . . . 85

3.1 Descriptive Statistics of Main Sample . . . 92
3.2 Lifecycle Bias and the Generalized-Errors-in-Variables Model . . . 98
3.3 Lifecycle Estimators in the Swedish Data . . . 103
3.4 The Revised Lifecycle Estimator . . . 113
3.5 The Lifecycle Estimator with Few Income Observations . . . 114
A.1 Intergenerational Elasticity Literature . . . 120
A.2 Moments of Income Process in Guvenen (2009) . . . 121

List of Figures

1.1 Evolution of Number of ENEM Test Takers 30
1.2 Example of application of the AA law in the state of Bahia. 30
1.3 Placebo Tests including 2 pre-periods 31
1.4 Heterogeneity of Effect by Baseline Share of PS Enrollment 32
B.1 Heterogeneity of AA Treatment by Baseline Share 44

2.1 Private vs Public Schools 69
2.2 Example of the Quota Law for the State of Bahia 70
2.3 Distribution of Treatment Variable 71
2.4 Estimate of Treatment Effects between School Systems 72
2.5 Estimate of Movements from Private to Public Schools by Ethnicity 73
2.6 Estimate of Movements from Private to Public Schools by School of Origin Socioeconomic Status 74
2.7 Estimate of Movements from Private to Public Schools by School of Origin Quality Index 75

3.1 Components of the Income Process 94
3.2 Trends in Lambdas and IGE Estimates 99
3.3 Illustration of Potential Problems with Fixed Effect Estimators 104
3.4 Estimation of Trends in the IGE 105
3.5 Extrapolating from Observable Profiles 107
3.6 A standardized errors-in-variables model in Swedish Data 116
A.1 Components of the Income Process 119
A.2 Variance of Income by Group 121

Chapter 1

Affirmative Action, Centralized Admissions and Inequality in Access to Higher Education: Evidence from Brazil

1.1 Introduction

Inequality in access to higher education is a growing concern in both developed and developing economies. In the U.S., Haveman and Wilson (2007) find a gap of almost 50 percentage points in college attendance between students in the top and bottom quartile of family income. In Brazil, according to the population Census of 2010, the share of college enrollment for individuals aged 18 to 22 is equal to 3.7 percent in the lowest quartile and 34.2 percent in the top quartile of family per capita income. Higher education is an important indicator of future labor market outcomes.[1] Therefore, the barriers to college attendance faced by disadvantaged students contribute to the perpetuation of income inequality and to lower social mobility.[2]

In this context, a wide range of policies have been implemented to improve access of vulnerable groups to higher education in different countries. Examples of such policies include: the expansion of financial aid, the reduction of tuition costs, the enlargement of the public and technical postsecondary systems, the practice of affirmative action in elite colleges and the reduction of application costs. Following the spread of these interventions, a recent grow-

[1]Although highly heterogeneous, the literature has found consistent average positive returns to higher education (e.g. Rodriguez et al. 2016; Zimmerman 2014).

[2]According to Haveman and Wilson (2007), the increase in U.S. income inequality over the past three decades has increased inequality in educational attainment between those with high and low incomes. This creates a vicious circle which perpetuates inequality. In this line, Chetty et al. (2017) explore the role of colleges in the transmission of inequality across generations in the U.S. They find that the fraction of students from low-income families fell sharply at colleges with the highest rates of bottom-to-top-quintile mobility between 2000-2011.

ing body of the economic literature has investigate how such initiatives affect the structure of higher education systems in terms of equality and diversity of the student body.

In this paper, I add to the literature by analyzing how two large governmental interventions affected enrollments of disadvantaged individuals in elite public higher education institutions in a particularly interesting institutional context. On the one hand, Brazil is one of the most unequal countries in the world; returns to higher education are relatively high,[3] while access is extremely unequal, especially for the elite colleges. On the other hand, most of these top institutions are public. Moreover, governmental expenditure in higher education is substantial compared to the primary and secondary levels, reinforcing the cycle of inequality.[4]

To reduce this problem, the government of Brazil recently adopted a set of policies that aimed to expand enrollments of disadvantaged groups to public colleges. The most important ones were the implementation of a centralized admission system in 2010 - the *Sistema de Seleção Unificada* (SISU) - and of a national *Affirmative Action Policy* (AA) in 2013. The SISU centralized applications in a nationwide online platform available for undergraduate degrees of federal and state institutions, based exclusively on grades of a national standardized exam. By reducing application costs, SISU created a more efficient market (Machado and Szerman, 2018). Yet, little is known about how it affected equality in access to public higher education. The AA, in contrast, determined that fifty percent of all vacancies in undergraduate degrees at federal institutions have to be reserved to students that attended public schools during all 3 years of secondary education,[5] with a certain percentage destined to race and economic-based minorities. By prioritizing equality, AA does not only mechanically increase enrollments of students of low socioeconomic status (hereafter low-SES), but also impacts incentives differently for each demographic group. In spite the large magnitude of these reforms, to the best of my knowledge, no academic papers tried to uncover their distributional effects. I try to fill this gap by investigating: (i) how SISU and AA changed the demographic composition of elite higher education institutions; (ii) if and how the two policies interact; (iii) the mechanisms behind the results.

My identification strategy relies on the progressive adoption of both policies, which gen-

[3] Brazil is the 11^{th} most unequal country according to its Gini Index (World Development Indicators, 2019). Earnings of workers with a tertiary degree are 2.5 times higher than the ones of workers with upper-secondary education. The OECD average is 1.56 (OECD, 2017).

[4] Brazil spends 3.8 thousand USD annually per student in primary education, while the OECD average is 8.7. In contrast, Brazil spends 11.7 thousand USD per student in tertiary education, similarly to European countries such as Italy (11.5) and Spain (11.8). The OECD average is 16.1, due to countries with substantially higher average spending, such as the US (29.3) and the UK (24.5) (OECD, 2017).

[5] According to data from ENEM 2016 (a National Standardized Exam), among the top 10 percent of high-schools in Brazil, 88% are private. Since attending a better (private) school is correlated with a higher socioeconomic status, the AA policy targets public school students as a proxy for low-economic status.

erates cross-sectional and time variation. Under a common trend assumption, this allows for the identification of their causal effect. The main specification includes time and degree-university fixed effects, in addition to a degree-level control for the number of vacancies and a municipality trend. The main identifying assumption for causal interpretation of the estimates is that trends in the outcome variable for treated and control units are parallel in absence of treatment. However, the existence of time-varying unobservable characteristics that are correlated with the outcome could be a potential threat for identification. In order to provide suggestive evidence that trends between controls and treated units are parallel, I conduct a series of placebo experiments, including both one and two lead variables, and find that coefficients for trends in pre-periods are close to zero and insignificant.

Results show that the full adoption of AA – from zero to fifty percent of reserved vacancies – increases enrollments of public school students (PS), non-white public school students (PSNW) and low income public school students (PSLI) by, respectively, 9.9, 7.0 and 2.4 percentage points, an increase of 18, 29 and 34 percent for the average program. These average effects mask remarkable heterogeneity. The full adoption of AA increases participation of PS by up to 28 p.p. and of PSNW by 15 p.p. for degrees in the lowest decile of PS participation share in the baseline year. Full adoption of SISU acts in the opposite direction, decreasing enrollments of these groups by 3.8, 2.8 and 4.1 percentage points, representing a negative effect of 7, 12 and 59 percent for the average program. Finally, the interaction between both policies creates an additional effect that increases enrollments of all vulnerable groups.

I present, then, suggestive evidence to shed light on the mechanisms behind these results. The full adoption of SISU substantially increases the average number of applications, enrollments of out-of-state students and average grades of enrollees. This suggests that centralization creates a more efficient market, as shown by Machado and Szerman (2018). Although the centralized system increases competition across all degrees, I find that it displaces low-SES students in the least competitive programs only. I then show that this pattern can be explained by the heterogeneous effects of SISU on the composition of out-of-state students in different degrees. In the least competitive programs, SISU disproportionately increases enrollments of high-SES out-of-state students, displacing the low-SES ones. Instead, in the most competitive programs, SISU disproportionately increases enrollments of low-SES out-of-state students, leading to no crowding-out effect. This suggests that mobility constraints are an obstacle for low-SES students accepted in degrees with lower expected future returns, but not to high-SES students or to low-SES students accepted in the most prestigious degrees.

The expansion of AA also increases the average number of applications. In addition, it increases enrollments of public school students in programs in which their participation was as high as 67 percent in baseline. This suggests that the positive effects of AA on enrollments of low-SES individuals are probably not only driven by the mechanical change in the admission rule, but also by a change in application behavior.

Finally, I discuss two potential channels behind the positive effects observed in the interaction of the two policies: one driven by the availability of additional information and another by the extra protection in face of higher competition. The effect of AA is higher when SISU is large. The full adoption of a centralized system induces all interested applicants to join the SISU platform. Once in the system, the low-SES students benefit from the information channel. By providing low-SES students with additional information regarding their own grades and cutoff scores, the platform enables them to better target their applications, leading to more enrollments. In parallel, the effect of SISU is larger when AA is fully adopted. This creates an additional protection channel that prevents the crowding-out effect. Thus, the joint adoption of the policies leads to more enrollments of low-SES students, displacing both local and out-of-state high-SES students.

My paper broadly contributes to the empirical literature that studies the impact of interventions that aim to democratize access to higher education. To cite a few examples, this literature has studied the effect of expanding financial aid and reducing tuition costs (e.g. Solis 2017; Rojas et al. 2013; Deming and Dynarski 2010), the enlargement of access to community colleges or technical postsecondary education (e.g. Goldhaber and Peri 2007; Denning 2017) and the effect of initiatives that aim to change low-SES students' aspirations and perceived value of attending college (e.g. Kaufmann 2014; Oreopoulos and Dunn 2013; Jensen 2010). Specifically, my paper adds to the strands within this literature that investigate the impact of affirmative action policies in undergraduate education and the effect of reducing application costs on enrollments of low-SES students.

Most of the studies related to affirmative action in higher education focus on the US experience and investigate how the practice of a racial preference has impacted minorities' overall enrollment, graduation attainment, major choice and labor market outcomes.[6] More closely related to my paper are the contributions of Hinrichs (2012) and Backes (2012), who study how affirmative action affects college enrollments and the demographic composition of universities. The lasting effects of such policies on long-term labor market outcomes begin by their differential effects on enrollments. Therefore, understanding the short-term impact of affirmative action on matriculation of minority groups (and displacement of the majority), and consequently on the composition of students at undergraduate degrees, is an essential step to learning the extent of its effect. Both Hinrichs (2012) and Backes (2012) explore the variation across time and state on U.S. colleges and find that the ban of affirmative action has

[6] Arcidiacono and Lovenheim (2016) and Arcidiacono et al. (2015) offer a comprehensive updated review of this literature.

no impact for the typical student and the typical college, but decreases under-representative minority enrollment at selective colleges. Although much can be learned from the evidence of affirmative action in the U.S., the structure of such policies in the Brazilian undergraduate higher education is more similar to the one in India, where institutions have precise quotas and admissions depend solely on grades.[7] Bertrand et al. (2010) and Bagde et al. (2016) study the effect of caste-based quotas in higher education in engineering colleges in India, and find that affirmative action substantially increases the probability of admission and attendance of disadvantaged casters.

Previous evidence on affirmative action in Brazil focused on the localized adoption of policies by specific institutions. Francis and Tannuri-Pianto (2012) find that the racial quotas introduced by the University of Brasilia increased the proportion of black students and of students from a lower socioeconomic background. Francis and Tannuri-Pianto (2018) conclude, then, that this policy increased the average years of education, college completion and labor earnings of the targeted group. Estevan et al. (2019) show that racial preferences in admissions at the University of Campinas leads toward a shift of admission to students of families of lower socioeconomic status. These papers find little evidence of behavioral effects on pre-college human capital accumulation.

I expand on this literature in different dimensions. First, I study the impact of affirmative action on the demographic compositions of universities in an institutional context that is remarkably different from both the United States and India.[8] Second, I analyze the effect of an affirmative action reform that was both nationwide, as it affected all federal universities of Brazil (17 percent of all undergraduate enrollment) and highly advertised and debated by the public, which differs from most of the experiences in the U.S., India or Brazil, in which affirmative action policies were analyzed in localized (mostly state-level) contexts.

[7]In the U.S., multiple factors play a role in college admissions, AA is expressed in form of a "racial preference" that is not precisely quantifiable and depends on specific college policies. In Brazil and India, institutions have precise quotas and admissions depend solely on grades. However, the demographic structure of the Brazilian population is more comparable to the US. Both countries have a diverse mixed-race population composed by descendants of Indigenous, Africans and European immigrants. Also, both the U.S. and Brazil endured a long period of slavery, which is historically one of the main sources of the current racial and social inequality in both countries.

[8]Brazil and India are among the most unequal countries in the world nowadays. According to Assouad et al. (2018), in both countries, the top 10% income share is greater than 50% of total pre-tax national income, compared to around 45% in the United States, and less than 40% in Western Europe. Also, while the middle 40% income share comprises 40% of income in the US, this share falls to 30% in Brazil and India. The sources and dynamics of inequality are, however, very different in both countries. In Brazil, the legacy of racial inequality from the almost 400 years of slavery is still very persistent. In India, inequality has its roots in the caste system. Moreover, while inequality in Brazil has been at stable high levels in the past 15 years, in India, it has persistently grown since 1980. Furthermore, AA in higher education is a relatively new policy in Brazil (starting in the early 2000s and expanding nationally in 2013). In India, the Constitution of 1950 already mandates affirmative action for lower-castes (Bagde et al., 2016).

The literature that studies the impact of application costs on students' decisions to apply and enroll in higher education has shown that more informed students attend higher quality colleges and that financial constraints and information are important drivers of mismatch (e.g. Dillon and Smith 2017; Hoxby and Turner 2015). Although application costs affect decisions of all students, the low-income individuals are usually the ones who benefit more from policies that reduce these costs. Bettinger et al. (2012) find that low-income individuals who receive assistance to complete the application for federal aid increase their likelihood of college assistance and persistence. Similarly, Carrell and Sacerdote (2017) show that providing high-school students at the margin of failing to apply with a college mentoring application program and application waivers increases their matriculation in higher education, especially for students at disadvantaged schools. Hoxby and Turner (2015) find that providing low-income high ability students with semi-customized information on applications and application waivers leads to more applications and admissions, especially to better-quality colleges. Bulman (2015) presents evidence that variation in access to SAT testing centers or mandatory in-school administration policies improves college-going, especially for high-ability students at the margin of attendance. Finally, Pallais (2015) shows that even a small reduction in monetary costs is capable of widening students' applications, leading the low-income ones to attend more selective colleges. In sum, this set of studies finds that relatively small and cheap interventions that reduce application costs can impact important decisions of low-income students and increase their investments in higher education.

I add to this literature by investigating how a much larger change in application costs impacts low-SES students' enrollments. In my context, the introduction of SISU represents an expressive reduction in applications costs in several dimensions: financial constraints, effort provision and information. The previous literature, instead, has focused on smaller-scale interventions: application assistance, financial waivers or provision of information to low-SES students. Thus, it is not clear whether the reduction of applications costs caused by the implementation of a centralized assignment should have similar effects as the policies previously studied. One key difference is that the introduction of a centralized system represents a large decline in costs for *all students*, not only the low-SES group. Thus, understanding its heterogeneous impacts on students' behavioral responses is essential for uncovering its distributional effects.

These are very relevant questions, both for the economic literature and for policy makers. Centralized assignments are largely adopted in different countries and educational levels. Plenty is known, especially theoretically, about its effects on improving allocation and coordination in different contexts. The empirical evidence, though, is still scarce, especially in the field of education. The paper by Machado and Szerman (2018), which explores the gradual implementation of SISU, is the first to provide empirical evidence on how the introduction of a centralized admission system improves allocation in an educational market. They find a positive impact on average grades of admitted students of institutions and programs that adopted the centralized system, interpreting it as an increase in the quality of matching between college and student. They also find that the adoption of SISU increases migration to attend higher education and drop-out rates after a year of enrollment. Yet, they do not explore the heterogeneous and distributional impacts of the introduction of centralized assignments and my paper fills this gap.

Additionally, the unique institutional context analyzed in my paper creates an opportunity to test whether the adoption of centralized admissions and of affirmative action quotas simultaneously creates interactions that change the effects that each would have had alone. This provides new evidence for both above-mentioned strands of the literature. Finally, I add to the literature in educational policy in Brazil and other middle-income countries, by quantifying how these interventions contributed to the changes observed in the student body in the recent period.

The remainder of this paper is divided as following. In section 2, I present an overview of the Brazilian higher education system and the institutional context of the reforms. In section 3, I describe the data sources and my main sample of analysis. In section 4, I explain the identification strategy and provide evidence in its support. In section 5, I show the main results, while in section 6, I discuss the potential mechanisms behind them. Finally, in section 7, I discuss the main issues regarding the internal validity of my empirical strategy and, in section 8, I present some policy implications of my findings.

1.2 Motivation

1.2.1 Brazilian Higher Education System

According to the Census of Higher Education of 2014, the Brazilian Higher Education System is comprised of 2368 institutions, 298 public and 2070 private. The public system is a mix of Federal (107), State (118) and Municipal (73) institutions, which correspond, respectively, to 17, 9 and 2 percent of the total undergraduate enrollment of around 6.5 million students.[9] Federal and State institutions are, by law, free of any charge, while Municipal schools usually charge some tuition. Private institutions, on the other hand, are a mix of profit (998) and non-profit (1172) organizations and although submitted to federal regulations and education standards, have complete independence regarding tuition fees and administration.

[9] Considering only undergraduate on-campus programs and students with an active enrollment status.

Public institutions (especially Federal) are widely recognized in the country by their average superior quality. For instance, the Federal institutions scored, on average, 3.6 in a scale of 0 to 5 of the *Índice Geral de Cursos 2014* (IGC), a quality index elaborated by the Ministry of Education based on performance evaluations of undergraduate and graduate programs. State institutions scored 2.8, and private institutions 2.6. Furthermore, among the universities only, 24 out of the best 25 are public, being 5 of the State administration and the other 19 Federal. On an alternative ranking - *Ranking Universitário Folha 2014*, elaborated by *Folha de São Paulo*, the newspaper of the highest circulation in Brazil - a similar pattern appears. Among the top 25 universities, 17 are Federal, 6 are State and 2 are private. Therefore, due to their high quality and free tuition, public institutions usually attract a large number of applicants. In spite of the high demand for their undergraduate degrees, competition is highly heterogeneous by institution and field of study.

Before 2010, the admission system for higher education was decentralized. Institutions organized their own admission exams, the so-called *Vestibular*, some months prior to the beginning of the academic year or semester. Once the application period was open, students had to opt for a particular degree and pay an inscription fee that would allow them to take the admission test of that specific program-institution. Furthermore, they had to be present on a certain date, time and location and exams were usually offered only in the city of the school. Exams were institution-specific and had a variety of structures and content that often required specific preparation. Because of both high costs of application and moving, the higher education market in Brazil used to be highly localized.

1.2.2 Centralization of Admissions

For the admission process of 2010, the Ministry of Education implemented a centralized admission system called SISU - *Sistema de Seleção Unificada*. This consists of an online platform, where universities offer their vacancies, and students, using only their grades in the national standardized exam, apply to the offered spots.

Only Federal and State institutions could adopt the SISU and they were free to choose if and how to do so. For instance, they could choose to adopt the system for all the available vacancies, partially or only for some degrees. Institutions adopted the system progressively. Although the adoption of SISU decreased costs of admission processes for both schools and students, many institutions were cautious about how the system would effectively work in the first years of implementation. Table 1.1 shows the gradual expansion of SISU adoption. While in 2010, the system was adopted by 55 institutions offering a total of 17 percent of all undergraduate vacancies in the Brazilian public higher education market, in 2015, the system reached 108 institutions and 59 percent of vacancies. It is also clear from Table 1.1

that the adoption of the centralized system is higher among Federal than State institutions. Since Federal institutions are funded by the national government, they are expected to follow the guidelines of the Ministry of Education. Indeed, by 2015, all but two federal institutions adopted SISU (fully or partially).[10] State institutions, on the other hand, had incentives to adopt the centralized system in form of additional resources to be used in student assistance.[11]

The SISU admission system is solely base on the score students obtained in the so-called ENEM - *Exame Nacional do Ensino Médio*. This is a national standardized exam handled by the Ministry of Education and available once a year across the whole country. The ENEM exam was created in 1998 as a means of testing the performance of high-school graduates. The importance of the exam has increased over time, as the number of takers went from 157.221, in 1998, to nearly 8.7 million, in 2015.[12] As shown in Figure 3.1 from Machado and Szerman (2018), the evolution of the number of test takers has two jumps. First, in 2004, the exam became a mandatory requisite for PROUNI, a government scholarship targeted at public school and low income students. Second, in 2010, the exam became mandatory for the SISU application.

In order to be used in admission processes, the Ministry of Education reformulated the exam in 2009. It was announced it would become more rigorous and more similar to the admission exams handled by public universities. It consists of 180 multiple-choice questions in four areas - Mathematics, Humanities, Sciences and Languages - and a written essay. Item Response Theory is used for grading, so the scores are comparable across years. Students had to pay a small fee (68 reais or about 20 dollars in 2016), which can be waved in case of financial need. High school graduates had incentives to take the ENEM exam before 2009, as its grade was used fully or partially in some admissions processes of private institutions and in a limited number of public institutions. There were also incentives for high-schools, as national rankings of institutions were elaborated by the Ministry of Education and spread nation-wide every year based on these grades. However, with the adoption of SISU in 2010, incentives increased, especially for students interested in pursuing higher education in public institutions.

Students take the ENEM exam around October or November of the year before starting higher education. The SISU opens its inscriptions in January and July and institutions can

[10]The two exceptions are UNIR and UFOPA.

[11]These additional resources could potentially be a confounder of the effect of SISU. Yet, I show that this is not the case. Although I include both Federal and State institutions in my main specification, results are robust when including only Federal institutions, which do not receive additional resources by adopting SISU. Due to the magnitude of the SISU reform, these additional resources in student assistance might play a marginal role in the enrollment decisions only.

[12]Since the number of test takers increase from 2010 to 2015, and some of the outcomes of analysis come from the ENEM dataset, this creates a potential selection problem, which I discuss in Section 7 of the paper.

choose if they want to offer all their spots in January or split their vacancies between the two yearly editions. Students pay no fee to participate in the system. The SISU is open for 4 or 5 days and students have to submit two choices of institution-major combination (program or degree) at the end of the application period. At the end of every day, the system updates cutoff scores of the programs depending on the people that have already applied. Based on that information, students are able to update their preferences and choose other options for which they might qualify. Only the two choices submitted when the system is closed count for their offers of admission. Candidates are accepted to their most preferred choice. In case they do not qualify, they become part of a waiting list.[13]

1.2.3 Affirmative Action

Access to public higher education in Brazil has historically been unequal. According to the Census of 2010, 85 percent of students aged 16 to 18 enrolled in high school attended a public institution,[14] while only 51 percent of incoming students in public higher education institutions are graduates from public high schools. Moreover, from the high-school enrollees aged 16 to 18, 47.5 percent were non-white and attended a public institution, while only 23 percent of first-year students in public universities were graduated from a public high school and also non-white. The lack of representation of economic and ethnic vulnerable groups in the public tertiary education created a social pressure for more equity in access to this governmental service.

The first public universities to implement affirmative action policies in their admission process were the State University of Rio de Janeiro (UERJ) and the State University of Bahia (UNEB) in 2003. Since then, many other public institutions adopted some type of affirmative action. However, there was not a national initiative that obliged them to do so and many of the country's best federal institutions were still resistant to adopt quotas. On

[13]This mechanism of applications is denoted *Iterative Deferred Acceptance Mechanism.* The specific case of Brazil has recently been studied theoretically by Bó and Hakimov (2016). They find that participants following the simple strategy of choosing the most preferred college in each period is a robust equilibrium that yields the Student Optimal Stable Matching. However, students may not always follow this strategy. Using data from SISU 2016, the authors show that from day 1 to 2, 5.8 percent of programs have their cutoff reduced. The iterative nature of the system might benefit students who are unaware of their true ability, as they are able to update their preferences and apply to programs they first believed to be out of reach. In spite of the positive "feedback feature", SISU is shown to produce some instabilities. For example, around 10 percent of programs have their cutoff reduced after the last update of preferences. This means that, for applicants with grades close to cutoff scores, there might be incentives for strategic applications. The authors conclude that the system could be improved with a combination of an iterative phase and second stage in which unmatched students would submit a list of preferences among the still available options.

[14]Note that the individuals of 16 to 18 years old that attend high school in Brazil belong to an already selective group. Around 30 percent of the youth aged 16 to 18 is out of education and other 28 percent is in age-distorted grades, i.e., still in primary or middle school.

August 29^{th} 2012, the Brazilian government approved Law 12.711/2012 - *the AA Expansion* -, stating that 50 percent of the spots in all majors in federal institutions should be reserved to individuals that studied the 3 years of their high school in public institutions.[15] The other 50 percent remained open to everyone. Among these reserved vacancies, 50 percent goes to students with family gross per capita income of 1.5 minimum wages or less. Furthermore, a minimum share of the reserved spots is destined to black, mixed and indigenous students (non-white), according to the percentage of these ethnic groups in the population of each state (defined according to the last available population Census). Figure 1.2 contains an example of how the law would be applied to a specific case. Since the state of Bahia has 76.8 percent of non-white, 76.8 percent of the reserved spots or 38.4 percent of total spots have to be reserved to them.

Federal institutions had 4 years to implement the law, i.e., they had to reserve a minimum of 12.5 percent of places in 2013, 25 percent in 2014, 37.5 percent in 2014, reaching 50 percent in 2016, the latest. However, they could adjust to the 50 percent share immediately. Table 1.2 shows the gradual increase in vacancies reserved to affirmative action in Brazilian public higher education institutions from 2010 to 2015. Although reserved spots for quotas increase in every year from 2010 to 2015, it is evident that the adoption of AA, from year 2013 onwards, is responsible for a clear jump in the percentage of reserved undergraduate vacancies offered by institutions. Other interesting aspect is that the increase in reserved spots is driven mostly by the shift in quotas for non-white public school students (ethnic), rather than the vacancies destined to public school students independent of race (non-ethnic).

1.3 Data

This paper uses data from different sources. First, it uses the Census of Higher Education (CES) from 2010 to 2015, which contains information on the universe of students enrolled in undergraduate degrees in all Brazilian institutions. The student level module includes data regarding which program (combination of institution and major) he or she is enrolled in, and demographic information about the student, such as gender, age, ethnicity, place of birth and type of school attended in high school. The CES also comprises modules with data about the institutions and programs, such as number of vacancies and degrees, resources, installations and information regarding professors.

[15] Note that only federal institutions were required, by law, to adopt the AA policy. State institutions are not required to do so, as they are not under the federal government administration. Although in my main specification I include both Federal and State institutions, in order to provide results comprehensive for the whole public higher education system, my results are robust when including only Federal institutions, which can be argued to be the most exogenous sample.

I restrict the CES 2010-2015 student-level sample to incoming (or first-year) students only. I also keep only individuals from undergraduate on-campus programs (excluding all online learning degrees) and students that are reported to have been selected through a regular selection process for first year-students (excluding transfers and special programs). Finally, I maintain in the sample only students from Federal institutions and State universities, excluding individuals that attend State centers and institutes.[16] My final sample of incoming students comprises 2,282,078 individuals, distributed along the 2010-2015 period, as specified by Table 1.3.

Second, I use data from the National Exam of High School (ENEM) Microdata from 2009 to 2014. This dataset contains information on test scores of all students that took the exam in each year with their demographic and socioeconomic characteristics. A unique individual identifier (social security number) is used to link information from the ENEM dataset of year *t-1* with the CES data of incoming students of year *t*. The matched sample comprises 1,829,037 individuals and the matching rate increases from 65 percent in 2010 to 89 percent in 2015. Students from the CES sample that are not found in the ENEM Microdata probably did not take the exam in the year before, either because they were not high school graduates or because they were selected to programs that did not have the ENEM exam as a requirement. The increasing number of matches reflects the growing adoption of the ENEM grade as a requirement for admission in public institutions. Both datasets are only available for researchers upon approval of research projects at the headquarters of the National Institute for Educational Studies and Research (INEP), of the Brazilian Ministry of Education, in Brasilia.

Table 1.3 contains descriptive statistics of the main variables used in the analysis of this paper. It shows an overall increase in the participation of public school, non-white and low-income individuals in the student body, while the share of out-of-state students remained stable. The dummy variables of whether individuals attended a public high school and their ethnicity are based on the information individuals provide in the ENEM questionnaire.[17] When the variables are missing, the information is complemented with the answer contained in the CES dataset.[18] The variable of low income is defined as whether individuals come

[16]These individuals correspond to only 5.7 percent of incoming students in public institutions. They are excluded because the information regarding the adoption of AA in these institutions is missing.

[17]This might raise concerns about the strategic manipulation of the variables in the self-reported data. First, students have no incentives to lie when filling the ENEM questionnaire, as the individual students' answers are not shared with institutions by any means and cannot be used in the admission decisions. Second, to qualify for AA, students need to have attended all three years of high school at a public secondary institution. This information is not easy to be manipulated, as it needs to be proven with school certificates. In the ENEM questionnaire, high school students need to specify which school they are graduating from.

[18]By adopting this procedure, I reduce considerably the number of missing values for these two characteristics. However, the missing values are still a problem, especially in early years. I use the ENEM dataset as

from families with total income of less than one minimum wage and it comes from the ENEM questionnaire.[19] The student is defined as from out-of-state if the reported state of residence at the ENEM questionnaire (during the final year of high school) is different from the one where the individual is enrolled in higher education, reported in the CES. The information on grades is an average of the four multiple-choice sections of the ENEM exam: Math, Science, Humanities and Languages. It was standardized to have mean zero and variance one based on all the test takers of that year. Finally, the variables for the number of spots and number of applications per program comes from the CES dataset.

Although the CES dataset contains information on the complete population of first-year students enrolled in Brazilian public higher education institutions, I do not observe individual outcomes for all incoming students, as shown in Table 1.3. For instance, in 2010, I observe the type of school attended in high school for 70 percent of the individuals. This share reaches 97 percent in 2015. This selection problem could be an issue for my empirical strategy if the existence of missing values is correlated with the implementation of the treatments. In Section 7, I characterize this problem with detail and show that although the introduction of SISU increases the availability of information, the results of my empirical model remain stable after a series of robustness tests.

In addition to those two administrative datasets, two smaller ones were also used. First, the data of SISU was provided by the Ministry of Education of Brazil (MEC). It contains information on the number of vacancies of each program and institution offered through the centralized admission system from 2010 to 2015. Second, information on the AA expansion was gathered by the author using public documentation of the admission processes, as well as information provided directly by the institutions. The resulting dataset contains the number of vacancies destined to each category of affirmative action of each public institution in the country, from 2010, before the federal law, to 2015, when the law was nearly completely implemented. Additional detailed information regarding data sources, sample restrictions and constructions of each variable are available in the Data Appendix.

the primary source, because studenst answer whether they attended a public high school during *all 3 years* of secondary, a requirement for qualifying to the AA law. In the CES, in contrast, the type of school refers to the the institution where students graduated from in high school.

[19] Note that the income criteria used in the AA law is different: family per capita income of less than 1.5 minimum wage. My definition of low-income here includes only individuals that are substantially poorer.

1.4 Empirical Methodology

In order to study the causal effects of the implementation of SISU and of the AA Expansion on enrollments of low-socioeconomic status students, I use the following baseline model:

$$Y_{iput} = \beta_1 SISU_{put} + \beta_2 AA_{ut} + \beta_3 SISU_{put} * AA_{ut} + \gamma X_{put} + \alpha_{pu} + \alpha_t + \alpha_m * t + \varepsilon_{iput}, \quad (1.1)$$

where Y_{iput} is the outcome of student i, of program p, institution u and time t. The treatment variable $SISU_{put}$ ranges from 0 to 1 and is defined as the percentage of spots of program p, institution u, at time t that is offered through the centralized system. The treatment variable AA_{ut} also ranges from 0 to 1 and defines the percentage of vacancies at institution u and time t reserved to AA policies. The variable acquires the value one when a share of fifty percent of quotas are adopted, i.e., when the national law is completely implemented. Both variables are demeaned. The inclusion of the interaction between both policies captures the additional effect on the outcome observed when the policies are adopted simultaneously. The vector X_{put} identifies program controls, which, in baseline, is only the number of vacancies by degree. Finally, I include program fixed effects α_{pu} (which also absorb institution fixed effects, since every program defines major-institution combinations), time-fixed effects α_t and a local linear trend $\alpha_m * t$, defined at the municipality level.[20]

The main identifying assumption for causal interpretation of parameters β_1, β_2 and β_3 is that dynamics in the outcome variable for treated and control units are equivalent in absence of the treatment. This is analogous to the parallel trends assumption required for exogeneity in a differences-in-differences framework. In this setting, however, the model has multiple periods, multiple treatment and control units and a continuous treatment variable. Yet, the identification assumption is similar. In this type of specification, the presence of program-institution fixed effects absorbs all unobservable time-invariant characteristics at program or institution that might be correlated with the outcome. However, the existence of time-varying unobservable characteristics that are correlated with the outcome could still be a threat to causal identification.

Since the treatment variables are continuous, any binary definition of control and treatment group would be arbitrary. Therefore, for illustrative purposes, I present in Table B.1 Panel A, regressions of the main outcomes and covariates used in the analysis in year 2010 on the total treatment jump from 2010 to 2015. Although most of the coefficients are insignificant, the point estimates suggest that there might exist some differences in levels between treated and control groups in the baseline. These differences are somewhat expected. In 2010, some institutions or programs had already adopted both of the policies and, therefore, are

[20]The subscript m is omitted in the other variables for the sake of simplicity.

expected to have different characteristics if compared to the ones that had not yet adopted them. This difference in levels does not invalidate the exogeneity assumption required to my specification, as long as trends in the outcomes do not change differently between treated and controls. As in a simple differences-in-differences framework, the parallel trends assumption cannot be directly tested. Nevertheless, I run a placebo experiment to provide suggestive evidence that, in pre-periods, this assumption is valid.

Table 1.4 shows results of this exercise for the main outcome variables analyzed in the paper: proportion of public school students (PS), non-white public school students (PSNW), low-income public school students (PSLI) and out-of-state students. As it is usual in the empirical literature for fixed effects models (e.g. see Autor 2003), in addition to the treatment variables $SISU_{pu,t}$ and $AA_{u,t}$, I include the lead variables $SISU_{pu,t+1}$ and $AA_{u,t+1}$, which capture the existence of possible trends for the outcomes. Unlike the coefficients for $SISU_{pu,t}$ and $AA_{u,t}$, the coefficients for the lead variables are close to zero in magnitude and insignificant, providing suggestive evidence that, despite differential levels, in absence of treatment, trends in the outcome variables between treated and control units would likely have been parallel. Additionally, Figure 1.3 presents evidence of an equivalent experiment using two pre-periods instead of one. Results are analogous and confirm that treated and control units had similar pre-trends for the outcomes also when one additional period is included in the placebo test. These exercises show that institutions do not adopt SISU or AA as a response to changes in the outcomes observed one or two years before the implementation of the policies. This means that the adoption of both policies is orthogonal to previous changes in the observed demographic characteristics of the student body. Unfortunately, due to the lack of additional data besides the years 2010 to 2015, used for analysis, the inclusion of additional periods could only occur at the expense of reducing the data sample for testing.[21] Yet, the absence of trends is confirmed by more flexible placebo tests, including either a state-linear trend or no linear trend at the local level, as shown in Table B.2.

1.5 Main Results

1.5.1 Enrollments of Low-SES students

Table 1.5 presents the main results of this paper. The adoption of SISU crowds out students from all vulnerable groups. A full adoption of SISU reduces the enrollments of PS, PSNW

[21]Note that for the placebo experiment with one lead variable, I use the data from period 2010 to 2014 to test the treatment implementation from years 2011 to 2015. For the experiment with two lead variables, I use data from 2010 to 2013 to test implementation from periods 2012 to 2015. Therefore, the lack of data from periods before 2010 prevents the testing of parallel trends with multiple pre-periods.

and PSLI by 3.8, 2.8 and 4.1 percentage points, respectively. This represents a decrease of 7, 12 and 59 percent for the average program, when compared to baseline shares. On the other hand, the expansion of AA increases the enrollment of these categories of students. More specifically, the full adoption of the national affirmative action policy (from zero to 50 percent of quotas) increases participation of PS, PSNW and PSLI by 9.9 p.p., 7.0 p.p. and 2.4 p.p, on average, an increase of 18, 29 and 34 percent, compared to baseline shares. The policies act in opposite directions. While SISU hurts the most vulnerable groups, AA increases their enrollments in the public higher education system. Interestingly, the full adoption of AA in the average program more than compensates the crowding out generated by SISU for the PS and PSNW groups. Yet, the effect of AA is comparatively lower for the PSLI. This possibly reflects the lack of effectiveness of the current income criteria used in the national AA law.[22] For example, 75 percent of the incoming students in federal institutions in 2012 belong to families with per capita income that fit the AA law target of 1.5 minimum wages or less. Therefore, it seems that this current income rule is excessively broad and does not benefit the poorer individuals among the public school students.

Finally, since the policies act in opposite directions, it is unclear if and how the simultaneous adoption of both would impact enrollments. The inclusion of a term of interaction between them in the main specification helps clarifying this question. According to Table 1.5, the simultaneous adoption of SISU and AA creates additional effects on enrollments that helps the vulnerable groups. For instance, the effect of SISU on PS enrollment at the mean level of adoption of AA is -3.75 p.p. At the full adoption of AA, it becomes approximately zero.[23] This means that, when AA is fully adopted, the interaction between the policies creates an additional effect, which eliminates the negative impact of SISU on enrollments of PS. Similarly, when SISU is fully adopted, the effect of AA becomes equal to approximately 14 p.p.[24] Therefore, when SISU is fully adopted, the effect of AA is 40 percent higher than when it is adopted at the mean.

1.5.2 Heterogeneity of Results

Table 1.5 presents the results on enrollments for the average program and university. Yet, since programs and institutions can be very different regarding their initial characteristics, results are expected to be highly heterogeneous. On Table 1.6 and Figures 1.4 and B.1, I analyze how some of these effects change depending on the initial share of enrollments of

[22] Remember that I define PSLI as individuals that attended a public school and belong to a family of total income of less than one minimum wage. In the AA allocation, the income criteria used is considerably broader, encompassing individuals of family income of less than 1.5 minimum wages per capita.

[23] Here, the effect of SISU is equal to β_1 + β_3 (AA-$\bar{AA}$)= -0.0375 + 0.0686 x (1-0.4).

[24] Here, the effect of AA is equal to β_2 + β_3 (SISU-$\bar{SISU}$)=0.0988 + 0.0686 x (1-0.4).

PS in the baseline year. Programs with less public school students in 2010 tend to be more competitive. Therefore, it is expected that AA affects their structure the most. Table 1.6 presents results for the baseline specification, with additional interactions of dummies for each quartile of initial share of PS and variables SISU, AA and SISUxAA.

The effect of AA on enrollments of group PS decreases as the initial baseline share increases. In the first quartile of PS participation at baseline year, the full adoption of AA increases enrollments by 23 p.p., in the second by 11 p.p., in the third by 4 p.p. and in the fourth it has a negative effect of 3 p.p. The adoption of SISU seems to have little effect in the first quartile and effects in the order of negative 3, 7 and 6 percentage points in the following quartiles. Finally, the effect of the interaction of the policies is positive, significant and of relevant magnitude in all four quartiles. Although the effect of AA is highly heterogeneous depending on the initial share of public school students, due to its mechanical component, the effect of the interaction between both policies does not seem to be related to this characteristic. Regardless the initial share of PS, the interaction of both policies creates an additional positive effect on enrollments that is not related to the mechanical effect that explains the heterogeneity observed for AA. A similar pattern is observed for the group of PSNW.

In Figure 1.4, I plot results analogous to Table 1.6, but instead of obtaining coefficients for different quartiles, I estimate coefficients for deciles of initial shares of PS enrollments and plot them against the median share of PS participation in the corresponding decile. The graph presents a pattern very similar to the one shown in Table 1.6. The effect of AA on enrollments of PS decreases remarkably as the baseline share increases. The effect of SISU also becomes more negative, while the impact of the interaction is relatively stable. Figure B.1 plots the coefficient of AA on both enrollments of PS and PSNW. Both are highly heterogeneous and decrease remarkably with the initial share of PS, although the coefficients for PSNW are flatter and do not turn negative in the later deciles.

1.6 Potential Mechanisms

The main results presented in Section 5 raise a number of questions related to the mechanisms behind them. Why centralizing admissions crowds out low-SES students from public higher education? It is clear that the AA policy increases enrollments for benefited groups and that this effect is highly heterogeneous depending on initial shares of enrollments of public school students. However, is the effect of AA purely mechanical or is there evidence of any effect on students' behavior? The interaction between AA and SISU is highly positive, of relevant magnitude and does not seem to be driven by the mechanical component of the effect of AA.

Why the interaction between centralizing admissions and expanding AA creates an additional enrollment-effect for vulnerable groups? Most of these questions cannot be answered with precision with this data and a reduced-form approach, as effects possibly include changes in application behavior and effort provision that induce general equilibrium and spillover effects. Yet, I will try to provide some additional suggestive evidence on the mechanisms behind them.

1.6.1 Why centralizing admissions crowds out low-SES students from the public higher education?

Results from Table 1.5 present clear evidence that centralizing admissions crowds out the low-SES students. This is not an obvious result. On the one hand, the reduction of application costs and better access to information benefit low-SES students disproportionately (Bettinger et al. 2012; Pallais 2015; Hoxby and Turner 2015). If this effect had prevailed, the adoption of SISU would likely have increased enrollments of low-SES students. On the other hand, theory predicts that centralization creates more competitive markets. This is confirmed in Table 1.7, which shows that SISU remarkably increases the average number of applications.[25] By lowering application costs and providing a nationwide online platform, SISU expands the market of public higher education as a whole. In this line, Machado and Szerman (2018) found that SISU improved the quality of matching between students and institutions, raising the average grades of admitted students and enrollments of students from out of state. Similar results are shown in Tables 1.7 and 1.8. A full adoption of SISU increases average grades of enrollees by 0.3 standard deviations and raises the participation of out-of-state students by 5 p.p. If this effect had prevailed, the implementation of SISU would likely have reduced enrollments of low-SES students, who, on average, have lower average grades than their high-SES counterparts.

In spite the sizeable displacement effects, the impact of SISU on enrollments of low-SES students is not uniform across the distribution of programs. Table 1.9 shows that the full adoption of SISU decreases enrollments of PS, PSNW and PSLI by 6.4, 5.2 and 7.9 p.p. in programs in the lowest quartile of competitiveness. In contrast, SISU has no crowding out effects in the top quartile. Meanwhile, SISU increases enrollments of out-of-state students in all quartiles of competitiveness. The effect is higher in the top quartile: 7.2 p.p. compared

[25] The average number of applications is defined at the major level and comprises the total number of applications from both SISU and the decentralized entrance mechanisms. In SISU, individuals can only apply to two programs. Therefore, one could expect that centralization would reduce the number of applications. Yet, SISU can increase the number of applications through two channels: increasing the number of applicants and creating the possibility of one additional application in the same institution. Before centralization, individuals could only choose one major in the same institution and, with SISU, they can now choose two.

to 5 p.p. in the bottom one. Taken together, these results suggest that the creation of a more efficient market alone is not sufficient to explain the heterogeneous pattern of the crowding-out effect.

In Table 1.10, I explore this matter further. I estimate the model separately for out-of-state private and public school students and find considerable differences. For the least competitive degrees (Quartile 1), SISU increases enrollments of out-of-state private school students by 6.6 p.p. and of out-of-state public school students by 4.3 p.p. This represents an increase of 372 and 114 percent respectively, if compared to baseline shares. On the other hand, for the most competitive degrees (Quartile 4), the effect is 7.5 p.p. and 6 p.p., an increase of 78 and 214 percent, respectively. Therefore, in Quartile 1, SISU increases enrollments of high-SES out-of-state students by 3.3 times more than of their low-SES counterparts. This is probably an important driver of the displacement effect observed. In contrast, in Quartile 4, SISU increases enrollments of out-of-state public school students by 2.7 times more than of the private school ones, with no overall effect on enrollments of low-SES students. This suggests that these *extra* out-of-state public school students displace local public school students with lower average grades.

This pattern is corroborated with results from Table 1.11. In Quartile 4, SISU increases average grades of private school students by 0.15 standard deviations and of public school students by 0.26. Meanwhile, it has no impact on the demographic composition of enrolled students. Thus, for the most competitive degrees, SISU is able to attract better students in general. Yet, it is able to attract even better low-SES students, increasing efficiency with no cost for equity. In contrast, for the least competitive degrees, SISU displaces low-SES students, increasing efficiency through enrollments of high-SES out-of-state students, but decreasing equity in access.

In sum, the crowding out effect observed is due to the creation of a more efficient nationwide market combined with the existence of mobility constraints that affect low-SES students only. For the most competitive and more prestigious degrees, SISU is able to attract better students from out of state of different socioeconomic backgrounds. Due to the higher expected future returns of these programs, low-SES students are especially willing to endure the mobility costs of attending college out of their home state. On the other hand, for the least prestigious degrees, costs of mobility become higher than expected returns of attending college out of state for the low-SES students only. Therefore, in this case, SISU attracts out-of-state students that are disproportionately from a higher socioeconomic background, crowding out the PS, PSNW and PSLI groups from public higher education.

1.6.2 Is the effect of AA on enrollments purely mechanical or is there evidence of any effect on students' behavior?

From Table 1.5, we learned that the full adoption of AA increases the average enrollments of PS students by 10 p.p. We have also learned that this effect is highly heterogeneous regarding the initial share of PS students. For example, in programs in the first decile, the median share of PS enrollment is 16 percent and the effect of AA is 28 p.p; while in the 8th decile, the median share of PS enrollment is 77 percent and the effect of AA is zero. From this, it is evident that the effect of AA has a large mechanical component. This means that, by construction, AA will have an effect that is larger the lower the initial share of PS is. However, it is still unclear if, in addition to the mechanical effect, there is also a behavioral component.

The behavioral effect comes from the changes in the application behavior and the composition of applicants. The change in incentives caused by AA could impact not only the number and the pool of applicants, but also their effort and their grades. Hence, the behavioral effect can be described as the impact of AA on enrollments after the policy had been announced, but before the change in admission rules resorts applicants according to their demographic group. The mechanical effect, on the other hand, is the result of the AA policy implementation, had the applicants' behavior stayed the same. Without individual application data, it is not possible to precisely quantify these two different effects. Yet, with my dataset, some indirect evidence can be drawn.

In Table 1.7, I investigate how AA impacts the average number of applications. This effect is, a priori, unclear. Although the introduction of AA can encourage benefited groups to apply, it can also discourage the groups that would not benefit from the policy. The net effect could be positive, negative or zero. According to Table 1.7, results show that the encouragement effect that AA exerts on the benefited groups more than outweigh possible negative effects it would have on other categories. This result shows that the implementation of AA impacts at least one dimension of behavior: the average number of applications. Nothing can be said about the composition of applications with the current data, but it is more likely that these additional applications come from the groups benefited from AA.[26]

Another question is whether changes in application behavior explain changes in the enrollment pattern. Results suggest that there is evidence of a behavioral effect of AA also on enrollments, although it is not possible to quantify this effect. In programs with a PS

[26] There is also the possibility of re-sorting in the applications. For example, private school students could not only apply less, but change their behavior and concentrate their applications in the least competitive programs. However, I do not find any evidence in support of this claim. The effect of AA on the average number of applications of the bottom quartile of competitiveness is zero. The positive effect found is stable and concentrated in the top three quartiles only.

share in baseline higher than 50 percent, had these shares remained constant in absence of treatment, a pure mechanical effect would not explain positive results on PS enrollments. If the announcement of AA had not impacted the composition and quality of applicants, AA would have zero effect for programs with 50 percent of PS students already. Yet, results from Figures 1.4 and B.1 and Table 1.6 show that this is not the case. For example, in the 5th, 6th and 7th deciles, the median share of PS students in the baseline is 51, 58 and 67 percent. Yet, a full introduction of AA still has a relevant positive effect of 9, 6 and 3 p.p. This can be due to behavioral changes such as more and better-targeted applications from PS students or an increase in effort provision from this group.

Finally, I investigate if there is evidence of any behavioral effect induced by AA on average enrollees' grades. Column 3 of Table 1.8 shows that the full adoption of AA decreases average grades of incoming students by 0.05 standard deviations. Yet, by including, as controls, dummies for PS, NW and LI, this effect vanishes. This suggests that the negative effect on enrollees' grades induced by the adoption of AA can be fully explained by the changes in composition. This result masks important between group heterogeneity. In Column 5, we see that the adoption of AA increases the average grades of private school students by 0.13 s.d., while it decreases average grades of public school students by 0.1 s.d. This is consistent with the fact the AA decreases enrollments of private school students - only the best ones of the group end up being accepted - and increases enrollments of public school students. In practice, this means that AA increases the variance of grades of the incoming cohort. While behavioral effects on grades induced by changes in effort and incentives are likely, this cannot be confirmed with the current dataset and empirical strategy.

1.6.3 Why does the interaction between SISU and AA creates an additional enrollment-effect for vulnerable groups?

According to Table 1.5, the simultaneous adoption of SISU and AA creates an additional effect on enrollments of PS, PSNW and PSLI students. For example, the full adoption of SISU reduces enrollments of PS students by 4 p.p., when AA is adopted at the mean. In contrast, its effect is negative 6.5 p.p., when AA is not adopted, and zero when AA is adopted at its maximum. Alternatively, while the full adoption of AA increases enrollments of PS by 10 p.p. at the mean of SISU, it increases enrollments by 14 p.p., when SISU is adopted fully, and by 7 p.p., when SISU is not adopted. According to Tables 1.7 and 1.8, there is no evidence that the interaction of the policies increase the average number of applications, the enrollment of out-of-state students or average grades.

Table 1.9 shows how the interaction effect changes across the distribution of program-competitiveness. In the least competitive programs (Quartile 1), the interaction is not statistically different from zero. In the top three quartiles, the interaction is higher in magnitude and significant. In Quartiles 2 and 3, SISU reduces enrollments of public school students by 5 and 3 p.p. at the mean of AA, while it has zero and a small positive effect (1.4 p.p.) when AA is adopted fully. From another angle, AA has a positive effect of 3 and 10 p.p. at the mean adoption of SISU. This effect reaches 8 and 14 p.p. if SISU is adopted fully. Finally, in the most competitive programs (Quartile 4), SISU has no crowding-out effect at the mean of AA, but increases enrollments of public school students by 4 p.p. at the maximum of AA. Meanwhile, AA has an effect of 23 p.p. at the mean of SISU and of 27 p.p. at its maximum.

What are the potential explanations behind this positive *complementary effect* between both policies? First, the adoption of SISU creates a national competitive market, displacing local public school students in the bottom three quartiles of competitiveness, as described earlier. Then, when AA is adopted at its maximum, it is able to further avoid that out-of-state students with better grades take additional vacancies from local public school students. AA creates, thus, an extra *protection channel*. This extra protection, for instance, reduces enrollments of out-of-state private school students in Quartile 4, as shown in Table 1.10 (interaction effect of negative 4.5 p.p.). Second, the adoption of SISU centralizes application data in a unique online easily accessible platform, creating the possibility of an *information channel*. On average, low-SES students have worse information regarding their own ability and the application procedure (Hoxby and Turner, 2015). Thus, the adoption of SISU, which provides better information on programs, institutions and their cutoff scores, is likely to benefit low-SES students the most. They can, then, better adapt their choices of major and institutions, targeting their applications according to program cutoffs. By applying more effectively, students that benefit from AA have an additional channel for increasing their enrollment. Further research needs to clarify the importance of each of these channels in explaining the results.

1.7 Robustness of Results and Extensions

1.7.1 Spillovers

The existence of multiple institutions and programs being treated simultaneously creates the possibility of spillover effects that might bias the baseline results. Spillovers occur when the outcomes of a certain unit of treatment are influenced not only by the changes observed in that specific unit, but also by changes in treatment of other units. This would be a violation of the *Stable Unit Treatment Value Assumption* (SUTVA), which requires that the potential outcome of one unit is unaffected by the particular assignment of treatment on other units. In the case of large reforms, such as the ones analyzed in this paper, the potential for spillovers occur in different dimensions. For example, factors such as location, type of major and quality of institution might be crucial in order to understand how treatment in one unit spills over to the others. In this section, I investigate the possibility of existence of local spillovers and spillovers across similar programs and institutions.

Since around 90 percent of students that attend public higher education in Brazil do it within-state, location is one of the most important determinants of college choice. First, I define a measure of treatment exposure at the municipality level. The idea is that treatment units are more likely to be impacted by changes in other units in the same municipality than by variations in institutions or programs outside of its locality. The constructed measures are defined as:

$$SpilloverSISU_{put} = \frac{VacanciesSISU_{mt} - VacanciesSISU_{put}}{TotalVacancies_{mt} - TotalVacancies_{put}}$$

and

$$SpilloverAA_{ut} = \frac{VacanciesAA_{mt} - VacanciesAA_{ut}}{TotalVacancies_{mt} - TotalVacancies_{ut}},$$

where m is the subscript for municipality.

Then, I run the baseline specification with these two additional measures of exposure. Table 1.12 shows that there are sizable local spillovers for AA. A full adoption of the AA policy by other schools in the same municipality of institution u decreases enrollments of PS and PSNW and PSLI at such institution by 6 and 3 p.p., respectively. This means that, when controlling by these negative spillovers, the estimates for the impact of AA on institution u itself are higher than in the baseline specification. The impact of a full adoption of AA at institution u increases from 9.9 to 13.0 for PS and from 7.0 to 8.7 for PSNW in comparison to the specification without the spillover exposure measure. Therefore, the baseline estimates for the AA policy may be downward biased and the estimates found in this section could be seen as an upper bound. In addition, in Table 1.7, we observed that AA did not have any impact on the share of students enrolled from out of state. Together, these results suggest that AA impacts enrollments within local educational markets mostly.

On the other hand, Table 1.12 shows no evidence of local spillovers for SISU, while Table 1.7 shows that SISU has a large effect on enrollments of students from out of state. Hence, in the case of SISU, treatment seems to be affecting individuals from outside the locality. If, in absence of treatment, the affected individuals would have attended a public institution within their state or municipality, this would consist of a violation of the STUVA. Although the size of the bias introduced in this case should not be large, since 90 percent of individuals attend public university within-state, the analysis of other dimensions of spillovers is important to rule out this possibility.

Therefore, I define a measure of spillover by field (or area of major) and by the level of competitiveness of the degree. I split programs in sixteen groups g, which define combinations between four different fields[27] and four different levels of competitiveness. The rationale is that treatment units are likely to be more impacted by programs within a similar field of knowledge and a similar level of prestige. The measures of exposure are defined analogously to the ones for local spillovers, just substituting the index m by the index g. Table 1.12 shows that the adoption of SISU by similar programs decrease enrollments of low-SES students at program p, while the adoption of AA at similar institutions increases enrollments at university u. In spite the relevancy of the spillover measures, the inclusion of these variables do not change the treatment effects of AA and SISU, nor the interaction. These results remain robust for additional measures of spillovers considered (Table B.3): local spillovers at the state level, spillovers across similar majors only and spillovers across similar majors and across institutions of a similar quality. The only variable that changes the magnitude of the baseline estimates of the treatment effects are the local spillovers measured at the municipality level for AA, minimizing our concerns regarding the violation of the STUVA.

1.7.2 Dynamics

In this section, I investigate the existence of dynamic effects in the estimates. SISU was adopted progressively and the size of the centralized system increased substantially from 2010 to 2015. Therefore, effects in the transition from a partially centralized system to a fully centralized system might change as individuals adjust to the expanding new mechanism of admissions. If low-SES students take longer to adapt and learn about SISU, its negative effect on their enrollments could fade with time. This is not what we observe in Table 1.13. Interestingly, the crowding-out effect generated by SISU is negative and relatively stable in all years. If anything, the negative effect of SISU on enrollments of low-SES students is more pronounced in 2014 and 2015 than in earlier years. This suggests that these negative effects are persistent and do not vanish naturally as individuals learn about the centralized system.

In contrast, the AA expansion occurred abruptly in 2013, with the introduction of the national AA law. In period 2010-2012, the expansion was restricted to minor changes in specific institutions. The national expansion of AA in 2013 was highly advertised and debated by the media, the politicians and the civil society. Therefore, it is not surprising that the effect

[27] i) Education; (ii) Humanities and Social Sciences; (iii) Sciences, Mathematics, Engineering and Computer Science; and (iv) Health, Biology, Veterinary and Agriculture.

of AA on enrollments of PS, PSNW and PSLI is higher in 2013-2015 than in period 2010-2012. Finally, we observe that the interaction between both policies is small and insignificant in period 2010-2012, while it is large and positive in period 2013-2015. Taken together, these findings suggest that the joint expansion of AA and SISU during period 2013-2015 was especially important for the sake of equality in access to higher education. It is likely that the adoption of the national AA law was determinant to bring low-SES students to the SISU system. Once in the system, these individuals benefit from the newly available information. Thus, the interaction between both policies improves access of low-SES groups remarkably in period 2013-2015.

1.7.3 Sample Selection

One of the most problematic internal validity issues of my empirical strategy concerns the selection of outcomes, as I do not have data on PS, PSNW and PSLI status for all the incoming students. This is so for two reasons. First, not all students take the ENEM exam in the year before admission. The share of admitted students found in the ENEM data increased from 65 percent in 2010 to 89 percent in 2015. Second, demographic characteristics of students are reported with a considerable number of missing values in years 2010 to 2012 for the CES data.[28] Information on PS status and ethnicity comes primarily from the ENEM dataset. When unavailable, they are complemented with the CES dataset. Information on low-income status is available in the ENEM data only.

First, I study whether the probability for the information to be missing is systematically correlated with treatment status. I estimate the main empirical model using, as the dependent variable, indicators that take the value 1 if information on these characteristics is available for individual i. A full adoption of SISU is correlated with an increase in the availability of information of the magnitude of 6 p.p. for PS and PSNW status, and of 10 p.p. for PSLI (Table 1.14). This is expected, since individuals are required to take the ENEM exam for applying to a SISU adopter-institution. As for AA, the relationship between treatment adoption and sample participation is weaker, if any at all.

Having established that the selection of the outcomes is not random with respect to treatment status, especially in case of SISU, I try to investigate whether this issue is biasing my results. Using data from baseline year, I compute the share of missing values of each outcome by program. Then, I create two restricted subsamples with programs in the bottom of the missing variables' distribution for each outcome. Sample 1 includes only programs in

[28]Some institutions failed to report these characteristics for the whole student body in these years. From year 2013 on, INEP increased efforts to check the quality of the data and reports from institutions increased substantially.

the bottom one half of the missing values' shares. The median program of this sample has missing values for 6, 8 and 14 percent of individuals enrolled in 2010 for the outcomes of PS, PSNW and PSLI respectively. Sample 2 includes programs in the bottom quartile and its median program has zero, 4 and 9 percent of missing values for the analyzed characteristics in the baseline year.

Then, I estimate the main empirical model in the restricted samples (Table 1.14). Presumably, in the samples with higher availability of information since the baseline year, the selection of outcomes will have little impact on the results. Estimates from Sample 1 are very similar to the ones in the full sample. Results from Sample 2, in contrast, are of a slightly smaller magnitude for the AA treatment and of similar magnitude for the SISU treatment. This suggests that the selection problem may be biasing my estimates upwardly. However, this bias does not seem to be substantial. Therefore, results in Sample 2 could be interpreted as a lower bound for my estimates of the effect of the policies.[29]

1.7.4 Additional Specifications

In order to check whether the results are robust, I run the analysis using different sample selections. First, I run the complete set of results restricting the sample for federal and state universities only. These institutions are relatively more comparable in terms of number of students, types and number of degrees offered, structure and quality. Moreover, state universities were not subject to the AA Expansion Law, which creates a comparison treatment group of relatively the same size as the control group. Results are very similar when compared to the whole sample. Second, I run results restricting the sample for federal institutions only and then for federal universities only. This sample cut is also justifiable, since only federal institutions are required to adopt the AA law. Again, results remain stable for different outcomes and specifications. These are shown in Table B.4.

Additionally, using my main sample of analysis, I run additional robustness tests, by deleting from the sample students that were not enrolled in the program by the end of the academic year (first-year dropouts) and by deleting students enrolled in programs with less than ten incoming students. Again, results remain stable for all outcomes and specifications. This means that the estimates are not driven by short-term effects of temporary movers nor from big shifts in small programs. I also run my results using only years 2011-2015, to minimize concerns regarding the high number of missing values in year 2010. These results are shown in Table B.5.

[29]This is a temporary work-around. For the outcomes of PS and PSNW status, I will be able to obtain the information for the current missing values in my sample by linking my dataset with students records in high school, in future visits to INEP. Furthermore, I will investigate this issue further for the PSLI as well.

Moreover, I run a whole set of results by the program level. This means that instead of considering the outcome at the individual level and clustering standard errors at the institutional-level, I use the average outcome at the program-level instead. With the program-level regressions, I obtain results using different program weights - unweighted, weighted by incoming students and weighted by log of incoming students (Table B.6). Additionally, I obtain results at the program level using different types of fixed effects and trends: state, municipality and institution linear trend and municipality-time fixed effects. Results are fairly stable and become even more precise in this last specification, as shown in Table B.7.

Finally, I run my analysis by institutional level, instead of program level (Table B.8). This means that I include my SISU treatment at institutional level and add institutional fixed effects, instead of program-institution fixed effects. Results are also stable to this analysis that explores only the more aggregate form of variation in the sample.

1.8 Concluding Remarks

In this paper, I analyzed how the adoption of SISU - a nationwide online platform that centralized admissions - and of AA - a policy that expanded affirmative action quotas to a uniform share of fifty percent in every major - impacted the demographic composition of Brazilian elite higher education institutions. I rely on cross-sectional and time variation in the implementation of such interventions to identify their causal effects. To confirm the exogeneity of my empirical strategy, I first show that pre-trends in the outcome variables are parallel between treated and control units, indicating that the adoption of the policies was not driven by changes observed in the student body in previous years.

There are some very special features in the context analyzed in my paper. First, I study a nationwide affirmative action initiative, implemented at elite higher education institutions, in a country with one of the highest levels of social and racial inequality in the world. Second, the adoption of SISU is a unique opportunity to investigate empirically the impact of a large reduction in different types of application barriers: financial, effort-related and information costs. Third, AA was adopted simultaneously with the expansion of SISU, allowing the investigation of the existence of complementarities between them. I find that SISU and AA act in opposite directions. While SISU crowds out students from public schools (PS), non-white from public schools (PSNW) and low-income from public schools (PSLI), AA increases their enrollments. Moreover, the interaction of both policies creates an additional effect that boosts enrollments of the benefited groups. Then, I show some possible mechanisms that explain the observed results.

Adoption of SISU increases the number of applications, average grades of admitted students and enrollments of students from out of state. Therefore, centralization creates a more efficient market, as it was previously shown by Machado and Szerman (2018). However, this hurts low-SES students financially constrained to move outside of their state. Moreover, students from a lower socioeconomic background are crowded out from degrees in the bottom half of the competitiveness distribution only. This suggests that mobility constraints are an obstacle for low-SES students accepted in programs with lower expected returns. Therefore, increasing governmental student assistance to reduce these financial constraints emerges as an important policy recommendation.

The adoption of AA, on the other hand, improves enrollments of PS, PSNW and PSLI. Although a large portion of this effect is mechanical and comes from programs that initially have shares of PS students lower than fifty percent - the minimum guaranteed by the AA law -, there is reliable evidence suggesting the existence of behavioral effects. First, AA increases the average number of applications. Second, enrollments increase even in programs with initial shares of PS as high as 67 percent. In such programs, AA is only expected to have an effect if it induces better applications of the targeted groups. Taken together, these results suggest that AA has an important role in improving access to public higher education. However, part of its effect could possibly be achieved by policies that enhance low-SES students' effort, confidence and beliefs, avoiding, thus, the stigma associated with AA.

I also find that the AA income criteria is too broad to encompass the poorer individuals among the public school students. Thus, the AA rules could target the lower-income individuals more specifically in order to increase their enrollments. Moreover, I show that AA increases the variance of grades of the incoming cohort. This has important policy implications. Higher inequality in the distribution of ability of the student body might have different implications for the learning environment, affecting the quality of teaching, interactions, peer effects, grading and, finally, dropouts, achievement and graduation. This is an important avenue for future research.

Finally, I discuss some potential mechanisms behind the positive effects found in the interaction between AA and SISU. Results show that the crowding out effect generated by SISU is persistent across all years, not fading away with time as individuals gradually adapt to the changes of the introduction of the new centralized system. In contrast, the introduction of the national AA policy in 2013 increases the magnitude of the AA effect and establishes the existence of a complementarity effect between both policies. Thus, AA seems to bring disadvantaged individuals to SISU, possibly by changing their beliefs regarding their suitability for higher education. Once in the SISU system, low-SES students likely gain the most from the availability of better information, as the centralized online platform provides easily accessible data on vacancies, programs, institutions and cutoffs. The second possible channel behind the interaction effect is more straight-forward. The simultaneous adoption of AA and SISU creates an additional form of protection, increasing the percentage of vacancies that are not available for the crowding out.

In conclusion, the policies acted in opposite directions. While SISU improved efficiency in detriment of equity, AA improved equity in detriment of efficiency. Therefore, the simultaneous adoption of these two interventions was particularly important to the achievement of both objectives. Furthermore, the simultaneous adoption of them created synergies that were particularly beneficial to lower-income students.

1.9 Figures and Tables

Figure 1.1: Evolution of Number of ENEM Test Takers

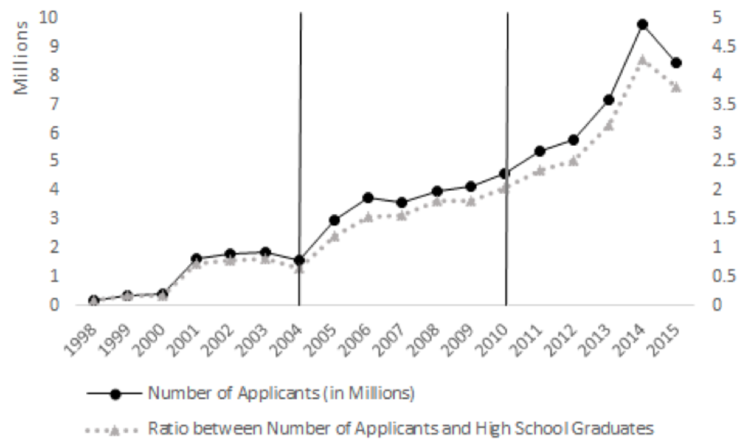

Notes: the number of ENEM test-takers increases from 157.221 in 1998 to 8.7 million students in 2015. In 2004, the exam became mandatory for the application to PROUNI, a government scholarship. In 2010, it becomes a pre-requisite for the SISU application. **Source:** Machado and Szerman (2018).

Figure 1.2: Example of application of the AA law in the state of Bahia.

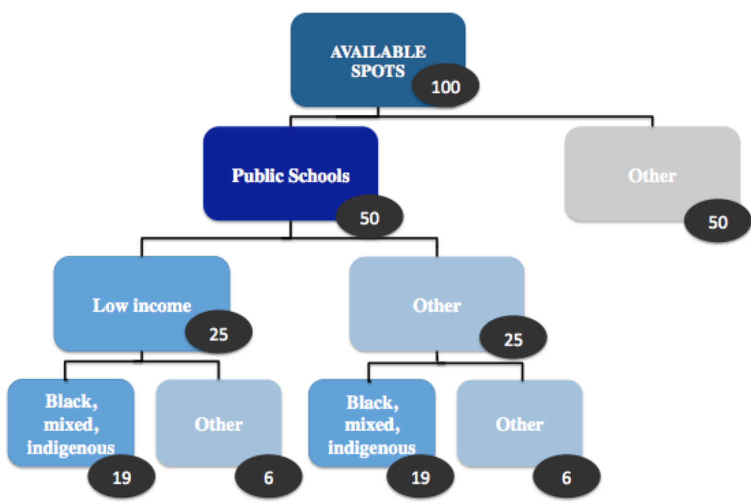

Notes: From every 100 vacancies offered by each major in a federal institution, 50 are reserved to students that attended all three years of secondary education in a public school, while the other 50 are open for regular competitive entrance. From the vacancies reserved to public school students, half are also reserved to individuals from families of per capita income of less than 1.5 minimum wage. Finally, the State of Bahia has 77 percent of non-white individuals according to the Census of 2010. Therefore, 77 percent of the reserved vacancies has to be destined to the non-white. **Source:** own elaboration based on Law 12.711/2012.

Figure 1.3: Placebo Tests including 2 pre-periods

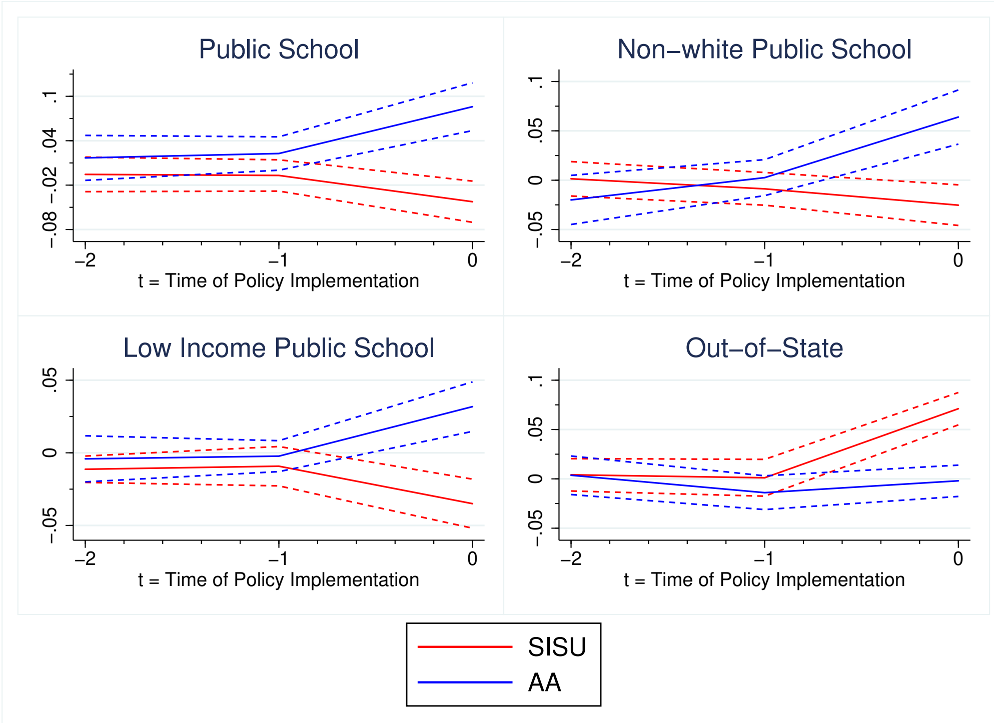

Notes: In this figure, I use the dataset from years 2010 to 2013 to test trends of two pre-periods for my main outcomes of interest: enrollments of public school, non-white public school, low-income public school and out-of-state students. In the graphs, I plot the estimated coefficients of $SISU_{pu,t}$ and $AA_{u,t}$, in time 0, and coefficients of the additional terms $SISU_{pu,t+1}$, $SISU_{put+2}$, $AA_{u,t+1}$ and $AA_{u,t+2}$, in times -1 and -2, respectively. Note that I use data from 2010 to 2013 to test whether the adoption of SISU and AA in periods 2012 to 2015 were correlated to changes in the outcomes observed one and two periods before implementation. The lack of data from years before 2010 prevents me from extending the analysis to further pre-periods. The estimation included time and program-institution fixed effects, a municipality linear trend and a control for the number of vacancies by program. Standard-errors are clustered at the institutional level and 95 confidence intervals are represented with dashed lines.

Figure 1.4: Heterogeneity of Effect by Baseline Share of PS Enrollment

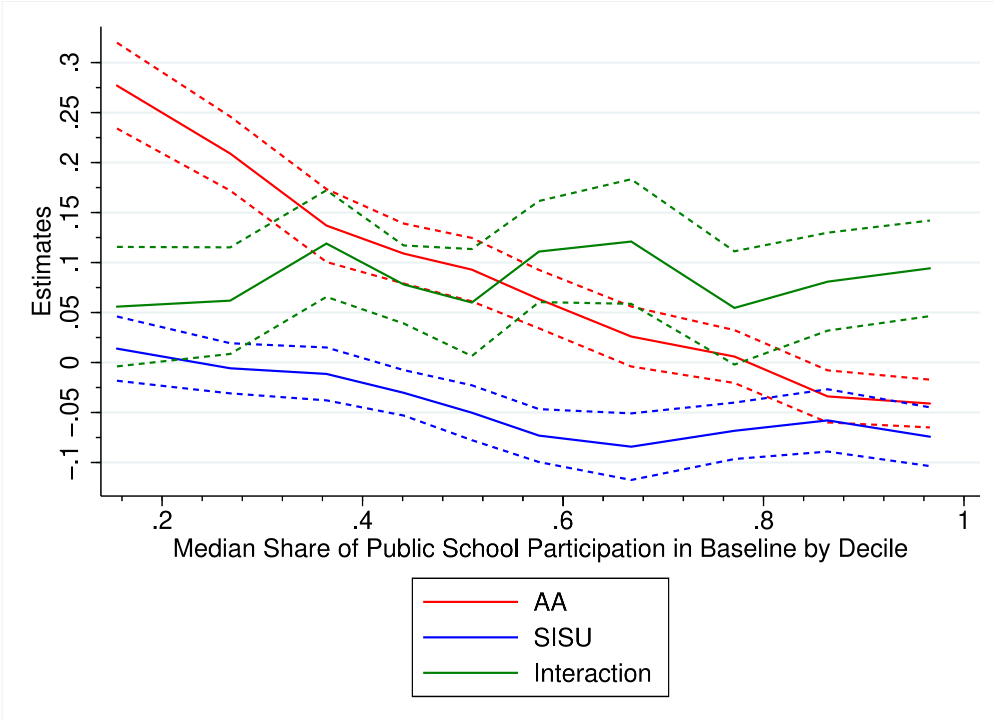

Notes: In this Figure, I plot the estimates of the treatment effects of AA, SISU and AAxSISU on enrollments of public school students (PS), estimated by decile of the initial share of PS matriculation in baseline. Using my main specification, I interacted dummies for each decile of PS in baseline with the treatment variables $SISU_{put}$, AA_{ut} - both demeaned -, and their interaction. The estimation included time and program-institution fixed effects, a municipality linear trend and a control for the number of vacancies by program. Standard-errors are clustered at the institutional level and 95 confidence intervals are represented with dashed lines.

Table 1.1: Expansion of SISU in Brazilian Undergraduate Education

	Public Institutions			Federal Institutions Only		
	Institutions	% Vacancies	% Programs	Institutions	% Vacancies	% Programs
2010	55	17.15	21.06	52	24.57	31.10
2011	73	27.57	37.06	68	37.93	51.68
2012	82	32.88	42.88	75	44.49	57.05
2013	88	40.54	50.35	77	50.25	63.27
2014	102	51.78	62.90	89	66.37	78.03
2015	108	58.95	72.24	92	72.54	86.28

Notes: This table shows the gradual adoption of SISU by Brazilian public institutions. In 2010, there are 131 public institutions, including both federal and state-owned. Out of these, 108 adopt SISU by 2015. Considering the federal institutions only, there are 94 in 2010, out of which 92 adopt SISU by 2015.
Source: SISU and Higher Education Census Datasets (2010-2015).

Table 1.2: Expansion of Affirmative Action in Brazilian Undergraduate Institutions

	Public Institutions				Federal Institutions Only			
	Institutions	%Total	%Ethnic	%Non-ethnic	Inst.	%Total	%Ethnic	%Non-ethnic
2010	69	19.93	8.46	10.87	42	17.73	7.53	9.74
2011	80	22.68	8.77	13.17	52	21.25	7.75	12.87
2012	82	24.54	9.20	14.60	54	23.43	8.44	14.33
2013	123	32.78	17.92	14.23	94	35.23	20.56	14.13
2014	124	38.95	20.83	17.36	94	42.34	24.52	17.18
2015	126	43.50	23.50	19.19	94	47.72	27.72	19.25

Notes: This table shows the gradual adoption of AA by Brazilian public institutions. In 2010, there are 131 public institutions, including both federal and state-owned. Out of these, 126 adopt AA by 2015. Considering the federal institutions only, there are 94 in 2010, all of which adopt AA by 2013, as required by federal law. Column *%Total* refers to the total percentage of vacancies reserved to AA quotas. Column *%Ethnic* refers to vacancies reserved to non-white and *%Non-ethnic* to vacancies reserved to public school students of every ethnicity.
Source: Documents of College Admission and Higher Education Census (2010-2015).

Table 1.3: Descriptive Statistics of Incoming Students in Public Higher Education Institutions

	2010	**2011**	**2012**	**2013**	**2014**	**2015**	**Total**
Public School	0.54	0.55	0.57	0.57	0.60	0.63	0.58
	(0.499)	*(0.498)*	*(0.496)*	*(0.494)*	*(0.490)*	*(0.483)*	*(0.494)*
% Non-missing	*69.8*	*88.3*	*89.3*	*97.1*	*96.8*	*97.4*	
Non-white PS	0.24	0.27	0.29	0.31	0.34	0.36	0.30
	(0.426)	*(0.442)*	*(0.452)*	*(0.461)*	*(0.474)*	*(0.480)*	*(0.460)*
% Non-missing	74.1	87.9	88.5	94.0	95.8	97.1	
Low-Income PS	0.07	0.12	0.10	0.13	0.15	0.16	0.12
	(0.255)	*(0.321)*	*(0.304)*	*(0.340)*	*(0.354)*	*(0.363)*	*(0.330)*
% Non-missing	*63.5*	*83.3*	*84.2*	*90.4*	*92.1*	*93.1*	
Out-of-State	0.10	0.09	0.09	0.10	0.10	0.10	0.10
	(0.301)	*(0.289)*	*(0.293)*	*(0.295)*	*(0.298)*	*(0.303)*	*(0.297)*
% Non-missing	*73.1*	*79.3*	*80.2*	*85.4*	*88.9*	*90.9*	
Grades	1.13	1.15	1.16	1.18	1.29	1.29	1.21
	(0.930)	*(0.879)*	*(0.891)*	*(0.948)*	*(1.013)*	*(0.961)*	*(0.943)*
% Non-missing	65.1	76.5	78.3	83.5	87.2	89.2	
N Applications	9.23	13.25	16.07	17.63	19.87	19.77	16.09
	(12.40)	*(16.74)*	*(25.32)*	*(22.25)*	*(23.65)*	*(22.76)*	*(21.46)*
Spots	91.53	97.87	95.36	98.52	95.60	92.93	95.31
	(121.53)	*(139.86)*	*(128.37)*	*(132.35)*	*(120.75)*	*(115.40)*	*(126.60)*
Observations	362634	370123	392865	383410	381464	391582	2282078

Notes: Values correspond to the mean. Standard deviations in parenthesis. *PS* stands for students that attended high school in a public institution, *Low-Income* for individuals from families of total income up to 1 minimum wage and *Out-of-State* for students that moved states to attend college. *Grades* is the average of the student standardized grade at the ENEM exam, *N Applications* is the average number of applications per program and *Spots* is the average number of vacancies per program. *Observations* are the total number of first-year students in my main sample of analysis.

Table 1.4: Placebo Experiment with Lead Variables

	PS	PS-NW	PS-LI	Out-of-State
SISUpu,t	-0.0406***	-0.0304***	-0.0383***	0.0554***
	(0.0113)	(0.0105)	(0.00782)	(0.00717)
SISUpu,t+1	-0.00142	-0.00394	-0.00602	-0.00341
	(0.00916)	(0.00686)	(0.00619)	(0.00889)
AAu,t	0.0944***	0.0777***	0.0293***	-0.00466
	(0.0165)	(0.0164)	(0.00740)	(0.00825)
AAu,t+1	0.0132	-0.0105	-0.00847	-0.00846
	(0.0106)	(0.00955)	(0.00520)	(0.00982)
SISUpu,t x AAu,t	0.0615***	0.0408*	0.0155	-0.0142
	(0.0222)	(0.0218)	(0.0123)	(0.0109)
N	1585361	1580230	1494373	1473111

Notes: Standard errors in parenthesis clustered at university level. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. Treatment variables are demeaned. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend.

Table 1.5: Effect of SISU and AA on Enrollments of Low-SES Students

	PS	PS	PS-NW	PS-NW	PS-LI	PS-LI
SISUput	-0.0421***	-0.0375***	-0.0319***	-0.0284***	-0.0427***	-0.0413***
	(0.0104)	(0.0103)	(0.00865)	(0.00903)	(0.00646)	(0.00662)
AAut	0.115***	0.0988***	0.0806***	0.0695***	0.0287***	0.0240***
	(0.0170)	(0.0151)	(0.0133)	(0.0131)	(0.00639)	(0.00601)
SISUput x AAut		0.0686***		0.0493***		0.0193*
		(0.0198)		(0.0176)		(0.0101)
Mean in baseline	0.54	0.54	0.24	0.24	0.07	0.07
N	2021455	2021455	2014838	2014838	1905968	1905968
R2-within	0.008	0.008	0.009	0.009	0.009	0.009

Notes: Standard errors in parenthesis clustered at university level. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income. Treatment variables are demeaned. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend.

Table 1.6: Effect of AA and SISU on Enrollments by Quartile of Baseline Share of Public-School Enrollment in Baseline

	PS	PSNW
Quart 1 x SISUput	-0.000926	-0.0101
	(0.0129)	(0.0120)
Quart 2 x SISUput	-0.0327***	-0.0311***
	(0.0119)	(0.0116)
Quart 3 x SISUput	-0.0744***	-0.0520***
	(0.0139)	(0.0131)
Quart 4 x SISUput	-0.0653***	-0.0323***
	(0.0130)	(0.0109)
Quart 1 x AAut	0.234***	0.132***
	(0.0185)	(0.0158)
Quart 2 x AAut	0.111***	0.0814***
	(0.0155)	(0.0140)
Quart 3 x AAut	0.0384***	0.0433***
	(0.0137)	(0.0156)
Quart 4 x AAut	-0.0309***	0.0121
	(0.0115)	(0.0167)
Quart 1 x SISUput x AAut	0.0545**	0.0598***
	(0.0256)	(0.0224)
Quart 2 x SISUput x AAut	0.0751***	0.0451**
	(0.0217)	(0.0205)
Quart 3 x SISUput x AAut	0.102***	0.0575**
	(0.0271)	(0.0240)
Quart 4 x SISUput x AAut	0.0844***	0.0483**
	(0.0225)	(0.0243)
N	1703287	1701226
R2-within	0.012	0.011

Notes: Standard errors in parenthesis clustered at university level. *PS* refers to students that attended high school in a public institution and *PSNW* to non-white public school students. Treatment variables are demeaned and interacted with dummy variables for each quartile of public-school enrollment shares in baseline. Quart 1 stands for the lowest baseline participation of PS. Controls for time and program-institution fixed effects, program number of spots and municipality linear trend.

Table 1.7: Effect on the Number of Applications and Enrollments of Out-of-State Students

	Log. Number of Applications		Out-of-State Students	
SISUput	1.053***	1.046***	0.0548***	0.0539***
	(0.118)	(0.119)	(0.00710)	(0.00666)
AAut	0.415***	0.431***	-0.0130**	-0.0104
	(0.130)	(0.133)	(0.00602)	(0.00659)
SISUput x AAut		-0.0771		-0.0105
		(0.162)		(0.0101)
Mean in baseline	9.23	9.23	0.10	0.10
N	2229570	2229570	1873289	1873289
R2-within	0.359	0.359	0.004	0.004

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *Out-of-state* refers to students that moved out of state to attend college. The number of applications is averaged by the number of vacancies. Results on number of applications are at the program level weighted by number of students. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend.

Table 1.8: Effect of AA and SISU on Average Grades of Incoming Students

SISUput	0.324***	0.292***	0.320***	0.291***	0.262***
	(0.0314)	(0.0276)	(0.0311)	(0.0274)	(0.0293)
AAut	-0.0630***	-0.00729	-0.0526**	-0.00387	0.129***
	(0.0236)	(0.0217)	(0.0249)	(0.0232)	(0.0247)
SISUput x AAut			-0.0415	-0.0136	-0.0172
			(0.0367)	(0.0334)	(0.0396)
SISUput x PS					0.0453**
					(0.0176)
AAut x PS					-0.229***
					(0.0208)
SISUput x AAut x PS					-0.00509
					(0.0393)
Mean in baseline	1.13	1.13	1.13	1.13	1.13
N	1807688	1753605	1807688	1753605	1753605
Controls for PS, NW, LI		Yes		Yes	Yes

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. PS, NW and LI refer to dummies that control for whether the individual comes from a public school, a non-white background and a low-income family. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Results are robust when grades include the essay component of the exam. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend.

Table 1.9: Heterogeneity of Effect by Quartile of Competitiveness of Degree

	Public School				**Non-white Pub.School**			
	Quart 1	Quart 2	Quart 3	Quart 4	Quart 1	Quart 2	Quart 3	Quart 4
SISUput	-0.0637***	-0.0497***	-0.0312**	0.00301	-0.0523***	-0.0373**	-0.0267**	0.00180
	(0.0158)	(0.0159)	(0.0127)	(0.0112)	(0.0149)	(0.0154)	(0.0117)	(0.00902)
AAut	0.0170	0.0300*	0.0977***	0.227***	0.0212	0.0417**	0.0780***	0.128***
	(0.0147)	(0.0177)	(0.0186)	(0.0262)	(0.0250)	(0.0174)	(0.0140)	(0.0175)
SISUput x AAut	0.0383	0.0869***	0.0746***	0.0631**	0.00822	0.0578**	0.0530***	0.0759***
	(0.0265)	(0.0278)	(0.0220)	(0.0306)	(0.0326)	(0.0250)	(0.0174)	(0.0209)
Mean in Baseline	0.80	0.63	0.47	0.31	0.40	0.30	0.20	0.10
N	368842	424804	430305	475323	365253	422655	430186	479352
R2-within	0.005	0.006	0.010	0.022	0.008	0.009	0.011	0.021
	Low Income Pub. School				**Out-of-State Students**			
	Quart 1	Quart 2	Quart 3	Quart 4	Quart 1	Quart 2	Quart 3	Quart 4
SISUput	-0.0793***	-0.0549***	-0.0241***	-0.00315	0.0499***	0.0445***	0.0421***	0.0721***
	(0.0111)	(0.00815)	(0.00609)	(0.00233)	(0.00770)	(0.00724)	(0.00938)	(0.0122)
AAut	0.0314**	0.0257***	0.0265***	0.0222***	-0.0141*	-0.00511	-0.0130	-0.00714
	(0.0129)	(0.00913)	(0.00592)	(0.00465)	(0.00786)	(0.00756)	(0.00806)	(0.0140)
SISUput x AAut	0.0322	0.0271*	0.0189*	0.0215***	-0.00157	-0.00221	-0.00682	-0.0294
	(0.0204)	(0.0145)	(0.0103)	(0.00752)	(0.00919)	(0.0118)	(0.0120)	(0.0206)
Mean in Baseline	0.19	0.08	0.04	0.01	0.06	0.09	0.12	0.12
N	332707	399052	414923	464728	330413	394992	410364	457728
R2-within	0.015	0.012	0.009	0.006	0.006	0.005	0.004	0.006

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Quart 1 stands for the least competitive quartile, while Quart 4 for the most competitive one. Competition is measured as the average grades of incoming students in the baseline. Results are robust when competition is measured as the average number of applications in baseline.

Table 1.10: Heterogeneity of Effect on Enrollment of Out-of-State Students by Quartile of Competitiveness of Degree and by Type of High School

	Quart 1	**Quart 2**	**Quart 3**	**Quart 4**	**All**
			Private School Students		
SISUput	0.0662***	0.0503***	0.0457***	0.0747***	0.0630***
	(0.0101)	(0.00946)	(0.00931)	(0.0141)	(0.00842)
AAut	-0.0201	-0.0164	-0.0127	-0.000234	-0.0125
	(0.0131)	(0.0110)	(0.0102)	(0.0138)	(0.00961)
SISUput x AAut	-0.00609	-0.0175	-0.0134	-0.0451**	-0.0230
	(0.0149)	(0.0168)	(0.0131)	(0.0226)	(0.0143)
Mean in Baseline	0.018	0.046	0.072	0.096	0.062
Growth	372%	109%	63%	78%	101%
N	64106	135866	195759	283840	772971
			Public School Students		
SISUput	0.0426***	0.0403***	0.0361***	0.0598***	0.0435***
	(0.00759)	(0.00723)	(0.0103)	(0.0110)	(0.00595)
AAut	-0.0116	0.00115	-0.00400	0.00205	-0.00350
	(0.00784)	(0.00620)	(0.0101)	(0.0166)	(0.00541)
SISUput x AAut	0.00696	0.00744	-0.00329	-0.00155	0.00380
	(0.00952)	(0.0112)	(0.0143)	(0.0229)	(0.00815)
Mean in Baseline	0.037	0.044	0.046	0.028	0.039
Growth	114%	91%	79%	214%	112%
N	257863	250234	204782	163385	1062281

Notes: In the first panel, results are estimated only for students that attended high school in a private institution, while in the second only for students that attended a public high school. Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Quart 1 stands for the least competitive quartile, while Quart 4 for the most competitive one. Competition is measured in average grades of incoming students in the baseline. Results are robust when competition is measured as the average number of applications in baseline. Growth is computed relative to the baseline share, after adding the respective point estimate.

Table 1.11: Heterogeneity of Effect on Grades by Quartile of Competitiveness of Degree

	Quart 1	Quart 2	Quart 3	Quart 4	Quart 1	Quart 2	Quart 3	Quart 4
SISUput	0.386***	0.415***	0.310***	0.179***	0.362***	0.358***	0.264***	0.148***
	(0.0390)	(0.0407)	(0.0404)	(0.0311)	(0.0359)	(0.0358)	(0.0371)	(0.0320)
AAut	-0.0389	0.00790	-0.0382	-0.181***	0.0594	0.139***	0.137***	0.0575**
	(0.0433)	(0.0366)	(0.0361)	(0.0286)	(0.0454)	(0.0378)	(0.0328)	(0.0278)
SISUput x AAut	-0.0440	-0.0687	-0.0449	0.0144	-0.0258	-0.0349	-0.0527	0.0292
	(0.0455)	(0.0599)	(0.0532)	(0.0446)	(0.0495)	(0.0635)	(0.0542)	(0.0461)
SISUput x PS					-0.0353*	0.0187	0.0535**	0.108***
					(0.0189)	(0.0223)	(0.0215)	(0.0312)
AAut x PS					-0.0968***	-0.175***	-0.240***	-0.383***
					(0.0249)	(0.0285)	(0.0278)	(0.0339)
SISUput x AAut x PS					0.0147	0.0163	0.0778	0.0856
					(0.0467)	(0.0523)	(0.0512)	(0.0647)
Controls for PS, NW, LI					Yes	Yes	Yes	Yes
N	312073	381241	398916	444560	304001	369096	384594	428469
R2-within	0.045	0.040	0.032	0.045	0.131	0.151	0.162	0.216

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. PS, NW and LI refer to dummies that control for whether the individual comes from a public school, a non-white background and a low-income family. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Quart 1 stands for the least competitive quartile, while Quart 4 for the most competitive one. Competition is measured as the average grades of incoming students in the baseline. Results are robust when competition is measured as the average number of applications in baseline.

Table 1.12: Spillover of Treatment Effects on Public Higher Education Market

	Public School			Public School Non-white		
SISUput	-0.0375***	-0.0315***	-0.0294***	-0.0284***	-0.0244***	-0.0234**
	(0.0103)	(0.0110)	(0.0109)	(0.00903)	(0.00909)	(0.00904)
AAut	0.0988***	0.130***	0.128***	0.0695***	0.0873***	0.0865***
	(0.0151)	(0.0167)	(0.0165)	(0.0131)	(0.0129)	(0.0128)
SISUput x AAut	0.0686***	0.0678***	0.0650***	0.0493***	0.0484***	0.0474***
	(0.0198)	(0.0190)	(0.0189)	(0.0176)	(0.0174)	(0.0174)
Spillover SISUput						
Municipality		-0.0151	-0.0171		-0.0113	-0.0123
		(0.0111)	(0.0110)		(0.0101)	(0.0101)
Field-Competition Level			-0.128***			-0.0711***
			(0.0193)			(0.0144)
Spillover AAut						
Municipality		-0.0610***	-0.0609***		-0.0317**	-0.0314**
		(0.0162)	(0.0160)		(0.0130)	(0.0129)
Field-Competition Level			0.127***			0.0207
			(0.0255)			(0.0218)
N	2021455	1991984	1991679	2014838	1985963	1985662

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Spillovers are defined at two levels: municipality and across similar major fields and competitiveness level of program.

Table 1.13: Yearly Effects

	PS	PSNW	PSLI	Out-of-State
2010xSISUput	-0.0528***	-0.0183	-0.0391***	0.0794***
	(0.0127)	(0.0113)	(0.00925)	(0.0121)
2011xSISUput	-0.0549***	-0.0348***	-0.0509***	0.0567***
	(0.0122)	(0.0103)	(0.00720)	(0.00800)
2012xSISUput	-0.0417***	-0.0287**	-0.0431***	0.0471***
	(0.0135)	(0.0116)	(0.00869)	(0.00799)
2013xSISUput	-0.0325***	-0.0273***	-0.0332***	0.0513***
	(0.0113)	(0.00973)	(0.00629)	(0.00932)
2014xSISUput	-0.0560***	-0.0432***	-0.0453***	0.0517***
	(0.0137)	(0.0109)	(0.00784)	(0.00693)
2015xSISUput	-0.0647***	-0.0607***	-0.0520***	0.0576***
	(0.0220)	(0.0153)	(0.0103)	(0.0176)
2010xAAut	0.0870***	0.0506**	0.0128	-0.00804
	(0.0213)	(0.0229)	(0.00824)	(0.00968)
2011xAAut	0.0943***	0.0580***	0.0242***	-0.00642
	(0.0161)	(0.0151)	(0.00707)	(0.00717)
2012xAAut	0.0879***	0.0639***	0.0207***	-0.0104
	(0.0151)	(0.0139)	(0.00679)	(0.00652)
2013xAAut	0.118***	0.0970***	0.0335***	-0.0108
	(0.0184)	(0.0141)	(0.00691)	(0.00802)
2014xAAut	0.137***	0.112***	0.0404***	-0.0186*
	(0.0232)	(0.0158)	(0.00845)	(0.0102)
2015xAAut	0.137***	0.130***	0.0456***	-0.0259
	(0.0352)	(0.0240)	(0.0117)	(0.0199)
2010xSISUputxAAut	0.0127	0.0293	0.0223	-0.0283
	(0.0342)	(0.0346)	(0.0152)	(0.0198)
2011xSISUputxAAut	0.0194	0.0235	-0.00398	0.00638
	(0.0242)	(0.0199)	(0.0124)	(0.0125)
2012xSISUputxAAut	0.0465*	0.0313	0.00813	0.00397
	(0.0251)	(0.0232)	(0.0131)	(0.0102)
2013xSISUputxAAut	0.0810***	0.0648***	0.0235*	-0.000429
	(0.0293)	(0.0212)	(0.0130)	(0.0206)
2014xSISUputxAAut	0.152***	0.0932***	0.0136	-0.00762
	(0.0412)	(0.0292)	(0.0185)	(0.0179)
2015xSISUputxAAut	0.165**	0.136***	0.0550*	-0.0228
	(0.0641)	(0.0432)	(0.0284)	(0.0461)
N	2021455	2014838	1905968	1873289
R2-within	0.008	0.009	0.009	0.004

Notes: Standard errors in parenthesis clustered at university level. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. Treatment variables are demeaned. Results include controls for time and program-institution fixed effects, program number of spots and a municipality linear trend.

Table 1.14: Treatment status and Sample selection

Panel A	Extent of Sample Selection			
	Obs$_{PS}$	Obs$_{PSNW}$	Obs$_{PSLI}$	Obs$_{Out}$
SISUput	0.0597***	0.0574***	0.0952***	0.115***
	(0.0227)	(0.0213)	(0.0211)	(0.0204)
AAut	-0.0423	-0.0251	-0.0354	-0.0258
	(0.0363)	(0.0237)	(0.0264)	(0.0215)
SISUput x AAut	-0.0495	-0.0461	-0.0208	-0.00941
	(0.0565)	(0.0452)	(0.0472)	(0.0353)
N	2238832	2238832	2238832	2238832

Panel B	Results in Full Sample			
	PS	PSNW	PSLI	Out-State
SISUput	-0.0375***	-0.0284***	-0.0413***	0.0539***
	(0.0103)	(0.00903)	(0.00662)	(0.00666)
AAut	0.0988***	0.0695***	0.0240***	-0.0104
	(0.0151)	(0.0131)	(0.00601)	(0.00659)
SISUput x AAut	0.0686***	0.0493***	0.0193*	-0.0105
	(0.0198)	(0.0176)	(0.0101)	(0.0101)
N	2021455	2014838	1905968	1873289

Panel C	Results in Sample 1			
	PS	PSNW	PSLI	Out-State
SISUput	-0.0303***	-0.0269**	-0.0367***	0.0461***
	(0.0103)	(0.0111)	(0.00770)	(0.00853)
AAut	0.0905***	0.0680***	0.0290***	-0.00159
	(0.0187)	(0.0150)	(0.00838)	(0.0135)
SISUput x AAut	0.0516**	0.0484**	0.00222	-0.0215
	(0.0223)	(0.0205)	(0.0148)	(0.0172)
N	954822	927841	901927	889247

Panel D	Results in Sample 2			
	PS	PSNW	PSLI	Out-State
SISUput	-0.0382***	-0.0351***	-0.0443***	0.0307***
	(0.0107)	(0.0115)	(0.0109)	(0.00678)
AAut	0.0506**	0.0529***	0.0454***	-0.00805
	(0.0234)	(0.0191)	(0.0127)	(0.0154)
SISUput x AAut	0.0606***	0.0418*	0.00171	0.00280
	(0.0222)	(0.0235)	(0.0216)	(0.0156)
N	478735	453525	434530	418006

Notes: In Panel A, Obs_{PS} is a dummy that takes value 1 if I have information on Public School Status for the student i enrolled, while it takes the value zero if the value is missing. Obs_{PSNW}, Obs_{PSLI} and Obs_{Out} are defined analogously for public school non-white, public-school low income and students that moved states to attend college. Sample 1 contains programs in the bottom one half of the distribution of missing values of the respective outcome in the baseline year. Sample 2 is defined analogously for the bottom quartile. Standard errors in parenthesis clustered at university level. Controls for time and program-institution fixed effects, program number of spots and municipality trend. Treatment variables are demeaned and standard errors are clustered at the university level.

Appendix

1.A Additional Figures and Tables

Figure B.1: Heterogeneity of AA Treatment by Baseline Share

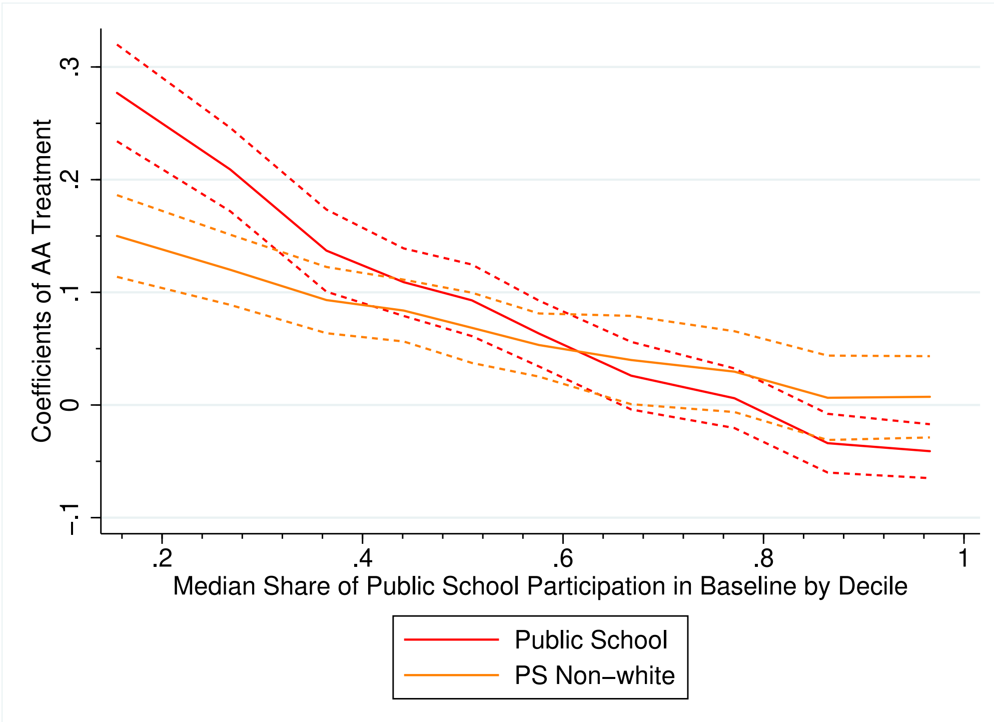

Notes: In this Figure, I plot the estimates of the treatment effects of AA on enrollments of public school students (PS) and non-white public school students (PSNW), estimated by decile of the initial share of PS matriculation baseline. Using my main specification, I interacted dummies for each decile of PS in baseline with the treatment variables $SISU_{put}$, AA_{ut} - both demeaned -, and their interaction. The estimation included time and program-institution fixed effects, a municipality linear trend and a control for the number of vacancies by program. Standard-errors are clustered at the institutional level and 95 confidence intervals are represented with dashed lines.

Table B.1: Treatment and Control Units

Panel a: Differences between Treated and Controls in Baseline

	AA Treatment		SISU Treatment	
	Coefficent	*Std. Error*	*Coefficent*	*Std. Error*
Public School	-0.0812*	(0.0437)	-0.0460	(0.0413)
Non-white PS	-0.0157	(0.0412)	0.0191	(0.0391)
Low Income PS	-0.0178	(0.0157)	0.0124	(0.0138)
Out-of-State	0.0306	(0.0261)	-0.0375*	(0.0211)
Grades ENEM	0.112	(0.119)	-0.224**	(0.110)
Non-white	0.0311	(0.0649)	0.110*	(0.0573)
Low Income	-0.0229	(0.0183)	0.0206	(0.0163)
Gender	-0.00898	(0.0182)	0.0179	(0.0135)
Age	-0.312	(0.388)	-0.284	(0.334)
Disability	-0.00133	(0.00357)	0.00247	(0.00264)
Spots	4.391	(19.18)	-2.493	(19.75)

Panel b: Distribution of 2010-2015 Jump in Treatment

	AA Treatment	SISU Treatment
Min	0.0	0.0
p(25)	0.0	0.0
p(50)	0.4	0.3
p(75)	0.8	1.0
Max	1.0	1.0
Mean	0.4	0.4
Std. Dev.	0.4	0.4
N individuals	362,634	309,404

Notes: In Panel 4a, I regress each covariate on a continuous variable equal to the total treatment jump from 2010 to 2015 of variables AA_{ut} and $SISU_{put}$, separately. Standard errors are clustered at the university level. In Table 4b, I present some descriptive statistics of the continuous treatment jump used in the regressions in 4a. *PS* stands for Public School and *Out-of-State* for students that moved states to attend college. The difference in the number of observations used in the regression of AA and SISU is due to programs that had enrollments in 2010, but not in 2015. These individuals are excluded from the SISU regression, in which treatment is defined at the program-level, but not from the AA one, in which treatment is defined at the institutional level.

Table B.2: Robustness of Placebo Experiment

	PS	PS-NW	PS-LI	Out-of-State
Panel A: State Linear Trend				
SISUpu,t	-0.0393***	-0.0264***	-0.0403***	0.0556***
	(0.0120)	(0.00963)	(0.00719)	(0.00693)
SISUpu,t+1	-0.00191	-0.00130	-0.00959	0.00279
	(0.0101)	(0.00691)	(0.00619)	(0.00816)
AAu,t	0.0794***	0.0647***	0.0249***	-0.00346
	(0.0170)	(0.0135)	(0.00725)	(0.00770)
AAu,t+1	-0.00875	-0.0243***	-0.0173***	-0.0118
	(0.0132)	(0.00895)	(0.00501)	(0.00849)
SISUpu,t x AAu,t	0.0476	0.0337	0.00951	-0.0207*
	(0.0306)	(0.0218)	(0.0114)	(0.0109)
Panel B: No Linear Trend				
SISUpu,t	-0.0396***	-0.0236***	-0.0351***	0.0523***
	(0.0124)	(0.00880)	(0.00622)	(0.00694)
SISUpu,t+1	-0.00854	0.00313	0.000162	0.00236
	(0.0113)	(0.00862)	(0.00616)	(0.00893)
AAu,t	0.0832***	0.0740***	0.0203**	0.00280
	(0.0189)	(0.0146)	(0.00873)	(0.00901)
AAu,t+1	-0.00602	-0.0163	-0.0105*	-0.0115
	(0.0138)	(0.0101)	(0.00588)	(0.00789)
SISUpu,t x AAu,t	0.0518*	0.0381	0.0262*	-0.0258**
	(0.0306)	(0.0246)	(0.0134)	(0.0110)
N	1585361	1580230	1494373	1473111

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. PS refers to public school, NW to non-white, LI to low-income and Out-of-State to students that used to live in a different state before enrolling in college. Results include controls for time and program-institution fixed effects and program number of spots. Panel A includes a state linear trend and Panel B does not include any linear trend. Results in Table 5 include a municipality linear trend.

Table B.3: Robustness of Spillover Measures

	Public School				
SISUput	-0.0315***	-0.0364***	-0.0313***	-0.0385***	-0.0371***
	(0.0110)	(0.0103)	(0.0101)	(0.0104)	(0.0108)
AAut	0.130***	0.0969***	0.0984***	0.0967***	0.0960***
	(0.0167)	(0.0151)	(0.0157)	(0.0151)	(0.0153)
SISUput x AAut	0.0678***	0.0661***	0.0683***	0.0670***	0.0676***
	(0.0190)	(0.0197)	(0.0195)	(0.0198)	(0.0195)
N	1991984	2021150	2021455	2021455	2021455
	Public School Non-White				
SISUput	-0.0244***	-0.0278***	-0.0217**	-0.0286***	-0.0271***
	(0.00909)	(0.00903)	(0.00839)	(0.00904)	(0.00923)
AAut	0.0873***	0.0690***	0.0646***	0.0688***	0.0677***
	(0.0129)	(0.0132)	(0.0121)	(0.0132)	(0.0129)
SISUput x AAut	0.0484***	0.0484***	0.0486***	0.0486***	0.0485***
	(0.0174)	(0.0176)	(0.0168)	(0.0176)	(0.0174)
N	1985963	2014537	2014838	2014838	2014838
	Public School Low-income				
SISUput	-0.0396***	-0.0415***	-0.0370***	-0.0407***	-0.0388***
	(0.00671)	(0.00661)	(0.00643)	(0.00657)	(0.00642)
AAut	0.0340***	0.0245***	0.0225***	0.0246***	0.0270***
	(0.00628)	(0.00607)	(0.00632)	(0.00594)	(0.00611)
SISUput x AAut	0.0181*	0.0198*	0.0190*	0.0191*	0.0195**
	(0.00981)	(0.0102)	(0.00998)	(0.00995)	(0.00980)
N	1879188	1905713	1905968	1905968	1905968
Municipality	Yes				
Field-Competition		Yes			
State			Yes		
Field				Yes	
Field-Quality					Yes

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. Controls for time and program-institution fixed effects, program number of spots and municipality trend. In each column, I test a different measure of spillover. Municipality and State refer to local spillovers at the respective levels. Field refers to spillovers measured at similar majors nationally; competition to programs in a similar level of competition and quality for institutions of a similar quality index. Estimators of Table 13 contain spillover measures at municipality and field-competition.

Table B.4: Robustness - Different Samples of Institutions

	PS	PSNW	PSLI	Out-of-State	Grades
			Panel A: Federal and State Universities		
SISUput	-0.0335***	-0.0269***	-0.0393***	0.0534***	0.300***
	(0.0115)	(0.0102)	(0.00715)	(0.00746)	(0.0342)
AAut	0.110***	0.0815***	0.0248***	-0.0106	-0.0708***
	(0.0190)	(0.0159)	(0.00609)	(0.00796)	(0.0253)
SISUput x AAut	0.0785***	0.0579***	0.00677	-0.00819	-0.000273
	(0.0223)	(0.0194)	(0.0115)	(0.0114)	(0.0404)
N	1811266	1805614	1715665	1684964	1626165
			Panel B: Only Federal Institutions		
SISUput	-0.0393***	-0.0330***	-0.0359***	0.0487***	0.299***
	(0.00990)	(0.00906)	(0.00657)	(0.00713)	(0.0322)
AAut	0.0885***	0.0684***	0.0165**	-0.00636	-0.0291
	(0.0143)	(0.0170)	(0.00754)	(0.00925)	(0.0338)
SISUput x AAut	0.0657***	0.0433**	0.0255**	-0.0176	-0.0804*
	(0.0187)	(0.0201)	(0.0116)	(0.0116)	(0.0417)
N	1478086	1477637	1423800	1410850	1377128
			Panel C: Only Federal Universities		
SISUput	-0.0368***	-0.0342***	-0.0320***	0.0463***	0.266***
	(0.0114)	(0.0105)	(0.00724)	(0.00832)	(0.0356)
AAut	0.0935***	0.0825***	0.0203**	-0.00651	-0.0597*
	(0.0196)	(0.0235)	(0.00849)	(0.0126)	(0.0350)
SISUput x AAut	0.0719***	0.0465*	0.0153	-0.0176	-0.0523
	(0.0221)	(0.0246)	(0.0135)	(0.0136)	(0.0451)
N	1267897	1268413	1233497	1222525	1195605

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Panel A includes Federal and State Universities, Panel B Federal Universities and Institutes, Panel C Federal Universities only, while baseline estimates contain all three type of institutions.

Table B.5: Robustness - Different Sample Selections

	PS	PSNW	PSLI	Out-of-State	Grades
	Panel A: No first-year drop-outs				
SISUput	-0.0349***	-0.0272***	-0.0395***	0.0493***	0.302***
	(0.0107)	(0.00925)	(0.00672)	(0.00611)	(0.0301)
AAut	0.104***	0.0716***	0.0254***	-0.0108*	-0.0644***
	(0.0157)	(0.0138)	(0.00601)	(0.00608)	(0.0243)
SISUput x AAut	0.0656***	0.0490***	0.0211**	-0.00774	-0.0439
	(0.0206)	(0.0187)	(0.0105)	(0.00914)	(0.0366)
N	1759779	1753761	1656371	1627609	1568136
	Panel B: Programs larger than 10				
SISUput	-0.0377***	-0.0285***	-0.0413***	0.0538***	0.319***
	(0.0103)	(0.00900)	(0.00662)	(0.00667)	(0.0316)
AAut	0.0988***	0.0697***	0.0240***	-0.0104	-0.0540**
	(0.0152)	(0.0131)	(0.00602)	(0.00659)	(0.0253)
SISUput x AAut	0.0682***	0.0488***	0.0192*	-0.0105	-0.0446
	(0.0199)	(0.0177)	(0.0101)	(0.0101)	(0.0369)
N	2017909	2011235	1902948	1870622	1805246
	Panel C: Drop 2010				
SISUput	-0.0468***	-0.0328***	-0.0477***	0.0522***	0.340***
	(0.00861)	(0.00762)	(0.00651)	(0.00677)	(0.0335)
AAut	0.0984***	0.0702***	0.0231***	-0.0110**	-0.0494**
	(0.0147)	(0.0104)	(0.00621)	(0.00558)	(0.0236)
SISUput x AAut	0.0845***	0.0548***	0.0272***	-0.00447	-0.00588
	(0.0189)	(0.0136)	(0.0102)	(0.0110)	(0.0400)
N	1773229	1751965	1679956	1613163	1575629

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. In Panel A, I delete all first-year dropouts; in Panel B, I keep only programs larger than 10 students and in Panel C, I keep only years 2011-2015.

Table B.6: Robustness - Results Collapsed at Program Level

	PS	PSNW	PSLI	Out-of-State	Grades
			Panel A: No Sample Weights		
SISUput	-0.0377***	-0.0270**	-0.0410***	0.0542***	0.340***
	(0.0116)	(0.0104)	(0.00770)	(0.00623)	(0.0322)
AAut	0.0807***	0.0517***	0.0183**	-0.0105*	-0.0403
	(0.0143)	(0.0131)	(0.00711)	(0.00577)	(0.0284)
SISUput x AAut	0.0742***	0.0592***	0.0243*	-0.00384	-0.0606
	(0.0231)	(0.0200)	(0.0123)	(0.0100)	(0.0438)
N	40000	40071	39894	39820	39673
			Panel B: Weight by number of students		
SISUput	-0.0366***	-0.0273**	-0.0406***	0.0540***	0.328***
	(0.0121)	(0.0112)	(0.00775)	(0.00723)	(0.0365)
AAut	0.0990***	0.0673***	0.0210***	-0.0115*	-0.0472*
	(0.0169)	(0.0160)	(0.00667)	(0.00677)	(0.0278)
SISUput x AAut	0.0666***	0.0532**	0.0242**	-0.00888	-0.0707
	(0.0234)	(0.0217)	(0.0118)	(0.0108)	(0.0441)
N	40000	40071	39894	39820	39673
			Panel C: Weight by log number of students		
SISUput	-0.0402***	-0.0289***	-0.0416***	0.0547***	0.338***
	(0.0110)	(0.0101)	(0.00762)	(0.00645)	(0.0334)
AAut	0.0868***	0.0584***	0.0178**	-0.0117**	-0.0398
	(0.0147)	(0.0135)	(0.00701)	(0.00590)	(0.0270)
SISUput x AAut	0.0680***	0.0535***	0.0238*	-0.00498	-0.0709*
	(0.0228)	(0.0197)	(0.0122)	(0.00956)	(0.0425)
N	39706	39774	39657	39613	39485

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Results at the Program Level are estimated after the computation of averages of each outcome by program and year. In each of the Panels, I test a different measure of program weight.

Table B.7: Robustness - Collapsed at Program Level and Different Local Controls

	UF Linear	Mun Linear	IES Linear	UF-Time FE	Mun-Time FE
			Public School		
SISUput	-0.0325**	-0.0366***	-0.0430***	-0.0261*	-0.0290**
	(0.0128)	(0.0121)	(0.0122)	(0.0134)	(0.0123)
AAut	0.0801***	0.0990***	0.0844***	0.0792***	0.127***
	(0.0174)	(0.0169)	(0.0179)	(0.0173)	(0.0185)
SISUput x AAut	0.0555*	0.0666***	0.0598**	0.0431	0.0617***
	(0.0300)	(0.0234)	(0.0274)	(0.0305)	(0.0203)
N	40000	40000	40000	40000	40000
			Public School Non-white		
SISUput	-0.0235**	-0.0273**	-0.0332***	-0.0150	-0.0186*
	(0.0106)	(0.0112)	(0.0114)	(0.00933)	(0.0112)
AAut	0.0556***	0.0673***	0.0779***	0.0470***	0.0717***
	(0.0131)	(0.0160)	(0.0222)	(0.0114)	(0.0141)
SISUput x AAut	0.0424*	0.0532**	0.0391	0.0460**	0.0651***
	(0.0216)	(0.0217)	(0.0273)	(0.0194)	(0.0172)
N	40071	40071	40071	40071	40071
			Public School Low-income		
SISUput	-0.0439***	-0.0406***	-0.0353***	-0.0399***	-0.0385***
	(0.00729)	(0.00775)	(0.00735)	(0.00684)	(0.00802)
AAut	0.0107	0.0210***	0.0238***	0.00775	0.0265***
	(0.00690)	(0.00667)	(0.00788)	(0.00737)	(0.00734)
SISUput x AAut	0.0190*	0.0242**	0.0129	0.0201*	0.0323***
	(0.0113)	(0.0118)	(0.0137)	(0.0115)	(0.0116)
N	39894	39894	39894	39894	39894

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. Controls for time and program-institution fixed effects and program number of spots. Results at the Program Level are estimated after the computation of averages of each outcome by program and year. In each column, I test a different type of local trend or local Fixed Effect. *UF* refers to State, *Mun* to Municipality, *IES* to institution and *FE* to fixed effects.

Table B.8: Robustness - SISU Treatment Collapsed at Institutional Level

	PS	PSNW	PSLI	Out-of-State	Grades
SISUut	-0.0410***	-0.0306**	-0.0487***	0.0533***	0.329***
	(0.0147)	(0.0125)	(0.00834)	(0.00822)	(0.0361)
AAut	0.0902***	0.0635***	0.0210***	-0.0110*	-0.0475
	(0.0163)	(0.0142)	(0.00658)	(0.00635)	(0.0340)
SISUput x AAut	0.101***	0.0658***	0.0271**	-0.0117	-0.104**
	(0.0262)	(0.0231)	(0.0112)	(0.0117)	(0.0481)
N	2054025	2048093	1932267	1896255	1829037
R2-within	0.093	0.132	0.097	0.075	0.267

Notes: Standard errors in parenthesis clustered at university level. Treatment variables are demeaned. *PS* refers to students that attended high school in a public institution, *NW* to non-white, *LI* to low-income and *Out-of-State* to students that used to live in a different state before enrolling in college. *Grades* is a composite measure of the different multiple-choice sections of the ENEM exam. Controls for time and program-institution fixed effects, program number of spots and a municipality linear trend. Here, the SISU treatment is collapsed at the level of the institution, instead of being included in the level of the program, as in the baseline estimate.

Chapter 2

Does Affirmative Action in Undergraduate Education Impact High Schools?

2.1 Introduction

Affirmative action (AA) policies in higher education have been adopted in different countries to mitigate inequality in access, performance and graduation. These initiatives are particularly important in countries that still suffer from a legacy of long-term institutions rooted in racial and economic inequality. As examples, we can cite India with the caste system, South Africa with the apartheid, and the United States and Brazil with slavery. There is a large and growing economic literature which investigates the effect of AA in undergraduate education in different dimensions of students' behavior and outcomes (e.g. the decision to apply to college, the choice of institution and major, performance, graduation and future labor market outcomes)[1]. In contrast, less is known about how these initiatives affect individual behavior *before* college. By creating heterogeneous incentive-schemes for different demographic groups, these policies might influence multiple dimensions of individual behavior even before university application. These might include, for example, students' human capital accumulation, choice of high school and social interactions. Assessing the impact of AA on students before college is not straightforward, since researchers do not know how middle and high school students are exposed to institutions they might be willing to attend in the future. However, it is vital for a complete understanding of the overall impact of AA and its welfare effects.

[1]See Arcidiacono and Lovenheim (2016) and Arcidiacono et al. (2015) for updated literature reviews on the topic.

My paper aims to contribute to this literature by analyzing the case of Brazil. On August 29^{th} 2012, the federal government approved the so-called "Quota Law" (hereafter *QL*). It established that 50 percent of vacancies in each major of each federal institution has to be reserved to students that attended all three years of secondary education in a public school.[2] Furthermore, it stipulates subquotas, within these 50 percent, to racial and economic-based minorities. Federal tertiary education institutions in Brazil are widely recognized in the country by their high-quality and are free of any charge. Therefore, competition for a spot is fierce.[3] By reserving a substantial percentage of vacancies to certain demographic groups, the federal government increased incentives for public school attendance in high school, especially for non-white and low-income individuals. In this context, the objectives of this paper are to understand and quantify whether *QL*: (i) increases mobility from private to public schools; (ii) has heterogeneous effects by ethnicity and type of school of origin; (iii) impacts the quality of public and private schools and, finally (iv) creates spillover effects for students in affected schools.

The adoption of *QL* creates a differential impact on each of the 107 federal institutions in Brazil, depending on their pre-reform levels of quotas. For example, some institutions already reserved as much as 50 percent of their vacancies to public school students (e.g. Federal University of Juiz de Fora - UFJF). Others, however, had no quotas whatsoever (e.g. Federal University of Pernambuco - UFPE). In sum, *QL* exogenously imposes that all institutions adjust their quota levels to 50 percent, creating cross-sectional variation across institutions and a continuous treatment variable. Although the effect of *QL* in each federal institution is clearly observed, I do not know how a student in middle school is exposed to treatment in each tertiary institution. To circumvent this limitation, I construct a measure of exposure to treatment for students in 9^{th} grade (the final year of primary education) based on factors that affect college choice, such as place of residence of the student and the location, size and quality of the institution. Details for the construction of this measure are presented in Section 3, but, essentially, higher education attendance in Brazil is highly localized. Approximately, 86 percent of enrollees attend college in the same municipality in which they resided before. This means that location of residence is a very important predictor of which tertiary federal institution a 9^{th} grader is exposed to.

Results show that a full exposure to *QL* increases mobility of 9^{th} graders from a private middle school to a public high school by 4.5 p.p. or 29 percent. This effect is stronger for non-whites, who benefit from an additional subquota, and to individuals that attended a low-socioeconomic status (SES) and low-quality private school. After determining that *QL* affects

[2] Public school attendance is used as a proxy of socioeconomic status. In Section 2, I give more details about public schools and how they compare to their private counterparts.

[3] Details on the Brazilian tertiary system also follows in Section 2.

public school attendance in the expected direction, I investigate a more intriguing question. Does this exogenous influx of ex-private school students impact the quality of public schools? If so, in which direction? Does it spill over to original public school enrollees? Alternatively, how this exogenous outflow of students impacts private schools? Does QL impact inequality between the private and the public high school systems? Preliminary results show that QL has a positive effect on public school quality. Even if the private school movers come disproportionately from the worse private schools, they still impact positively the quality of the public schools they move to. This effect is, however, concentrated in a small portion of public schools: the best ones. Although additional research is needed, results suggest that QL has a positive impact on public schools at the mean, but led to an increase in inequality across schools within the public system. In the next version of this manuscript, I plan to investigate whether QL generates spillover effects and if it impacts private schools and inequality between and within the public and the private systems.

These are questions of extreme importance, especially in countries similar to Brazil, where socioeconomic inequality is high and current and prospective resources for educational investments are limited. Brazil is the 11^{th} most unequal country in the world.[4] One of the main roots of this inequality is, precisely, the educational system. In contrast with most developed countries, Brazilian educational expenditures are concentrated in tertiary, in detriment of basic education.[5] Nevertheless, access to higher education (especially public) is extremely unequal and returns are substantial.[6] The adoption of QL aims, specifically, at mitigating this issue. However, most critics of this type of policy argue that inequality in access to higher education is generated much earlier in one's trajectory and that such initiatives may cause more distortions than solutions. This argument is partially correct. For example, in Brazil, public schools overall quality is considerably lower than the one their private counterparts (Figure 2.1). Moreover, around 15% of Brazilian students drop out of school right after 9^{th} grade. In this line, policies targeting the reduction of inequality in earlier cycles of education are extremely needed. On the other hand, AA in higher education might have overall effects that go well beyond reducing inequality in college access. If such policies are able to increase incentives for public school attendance, they might reduce inequality in the basic education

[4] Brazil is the 11th most unequal country according to its Gini Index (World Development Indicators, 2019).

[5] Brazil spends 3.8 thousand USD annually per student in primary education, while the OECD average is 8.7. In contrast, Brazil spends 11.7 thousand USD per student in tertiary education, similarly to European countries such as Italy (11.5) and Spain (11.8). The OECD average is 16.1, due to countries with substantially higher average spending, such as the US (29.3) and the UK (24.5) (OECD, 2017).

[6] According to the population Census of 2010, the share of college enrollment for individuals aged 18 to 22 is equal to 3.7 percent in the lowest quartile and 34.2 percent in the top quartile of family per capita income. In parallel, earnings of workers with a tertiary degree are 2.5 times higher than the ones of workers with upper-secondary education. The OECD average is 1.56 (OECD, 2017).

system, increasing its welfare effects. This is what I would like to ultimately understand with this paper.

My paper is specifically related to the literature that investigates how AA in undergraduate education impacts high school students' outcomes. One strand of this literature has focused on how such policies affect pre-college human capital investments. Bodoh-Creed and Hickman (2018) develop a structural model of college admissions framed as a contest in which the outcome is decided by the students' choice of human capital. Cotton et al. (2018) use a simple version of this model to derive testable predictions and find, through an experimental approach, that AA increases effort levels of the benefited group, while not affecting the non-benefited students. In contrast, using a differences-in-differences framework, Antonovics and Backes (2014) find no evidence that banning AA at the University of California affected human capital accumulation for high school students. There are some studies that look at a similar question in Brazil. Francis and Tannuri-Pianto (2012) and Estevan et al. (2019) find no behavioral effects of AA policies implemented in two different Brazilian universities on pre-college human capital accumulation. In contrast, Assunção and Ferman (2015) find that an AA policy implemented by the State University of Rio de Janeiro decreased investments by black students, the target group. The authors argue that their results might be explained by the aggressiveness of the specific AA policy they analyze, which differs substantially from the other two cases.

More related to this paper, a second strand of this literature investigates how AA in higher education affects secondary students' choice of school and high school integration. To the best of my knowledge, all these papers focus on the consequences of the so-called "Top-N percent plans" in the U.S., a type of AA policy that guarantees seats in State universities for all students ranked in the top N% of their high-schools. The contribution most related to mine is the one by Cullen et al. (2013), who analyze the case of Texas. They find that among the subset of students with both motive and opportunity for strategic high school choice, at least 5% enroll in a different high school to improve the chances of being in the top 10% and benefit from AA. Estevan et al. (2019) show that, although this pre-college arbitrage is detrimental for diversity in tertiary institutions, it increases high school integration. My paper adds to this literature in two key dimensions. First, it provides clear evidence of the existence of this arbitrage effect of AA in higher education in a context never studied before. In spite of some similarities in terms of overall objectives between *QL* and the "Top-N percent plans", the two AA initiatives are substantially different, in terms of policy-design and the institutional context in which they have been implemented. While the "Top-N percent plans" are localized - as they were implemented at the State-level in three U.S. States -, the *QL* affected all 107 federal institutions in Brazil, including some of the very best universities in the country.

Therefore, we might also expect larger results on the strategic school choice of individuals' and on how it sequentially affects high schools. Second, I study a policy that clearly creates incentives for public high school attendance. In the context of high socioeconomic inequality and, consequently, of high segregation of the public and private school systems, the Brazilian case creates an ideal scenario for testing the effectiveness of an AA policy that simultaneously improves equality in access to higher education and the quality of public schools.

The remaining of this paper is organized as following. In Section 2, I provide a background of the Brazilian educational system and of the implementation of QL. In Section 3, I describe the main datasets and how I constructed the measure of 9^{th} graders exposure to treatment. In Section 4, I present the empirical strategy and main results regarding the individual strategic school choice in response to QL. In Section 5, I show results related to how QL impacts quality of public schools. Finally, in Section 6, I present concluding remarks and discuss the future steps of this project.

2.2 Institutional Background

2.2.1 Brazilian Educational System

The basic mandatory education system in Brazil is comprised of 9 years of primary, followed by 3 years of secondary school.[7] Students start first grade at age 6 and should finish high school at ages 17 or 18, before entering university. Although the government offers universal access to all grades of basic education, the public system coexists with a large number of private schools. According to the Census of Basic Education of 2011, there are 63030 schools in Brazil that offer the second cycle of primary education (grade 5-9), being 80% public and 20% private. In secondary education, there are 26973 schools, 70% public and 30% private. Moreover, the private system encompasses 12.6% of total enrollments in the second cycle of primary and 12.2% in secondary. Apart from size, the public and private systems are very different in other key and more relevant dimensions. Private schools are, on average, of better quality than their public counterparts. From the top 100 high schools in Brazil, according to the National Standardized Exam of 2011 (ENEM), 93 are private. Moreover, from the 10077 schools evaluated, the 4799 private schools perform considerably better, as seen in Figure 2.1. Additionally, private schools' socioeconomic level is substantially higher, as also portrayed in Figure 2.1.

On the other hand, tertiary education in Brazil presents a contrasting scenario. According to the Census of Higher Education of 2014, the Brazilian Higher Education System was

[7] From 2013, pre-primary education for children aged 4 and 5 also became mandatory.

comprised of 2368 institutions, 298 public and 2070 private. The public system is a mix of Federal (107), State (118) and Municipal (73) institutions, which correspond, respectively, to 17, 9 and 2 percent of the total undergraduate enrollment of around 6.5 million students.[8] Federal and State institutions are, by law, free of any charge. Private institutions, in contrast, charge tuition fees that may vary substantially, but were, on average, equal to 898 Brazilian Reais per month in 2017 (or 95.8% of the minimum wage)[9]. Public tertiary institutions (especially Federal) are widely recognized in the country by their average high quality. For instance, the federal institutions scored, on average, 3.6 in a scale of 0 to 5 of the *Índice Geral de Cursos 2014* (IGC), a quality index elaborated by the Ministry of Education based on performance evaluations of undergraduate and graduate programs. State institutions scored 2.8, and private institutions 2.6. Furthermore, among the universities only, 24 out of the best 25 are public, being 5 of the State administration and the other 19 Federal. On an alternative ranking - *Ranking Universitário Folha 2014*, elaborated by *Folha de São Paulo*, the newspaper of the highest circulation in Brazil - a similar pattern appears. Among the top 25 universities, 17 are Federal, 6 are State and 2 are private. Therefore, due to their high quality and free tuition, public tertiary institutions usually attract a large number of applicants.

2.2.2 The Quota Law in Higher Education

Access to public undergraduate education in Brazil is highly competitive. For example, according to the Centralized Admission System of 2016 (SISU 2016), 2664001 students applied to 242864 vacancies in a Federal institution, a rate of 11 students per vacancy.[10] Therefore, only students with high grades are able to successfully obtain a spot in these competitive colleges.[11] As a consequence, access to public higher education in Brazil has historically been unequal. For instance, 85 percent of high schools students aged 16 to 18 go to a public school, while only 51 percent of incoming students in public higher education institutions are graduates from public high schools. Moreover, 47.5 percent of high school students are non-white and go to a public school, while only 23 percent of first-year students in public

[8] Considering only undergraduate on-campus programs and students with an active enrollment status.

[9] According to data from *Mapa do Ensino Superior no Brasil 2017*, from the *Sindicato das Mantenedoras de Ensino Superior*. Information available at:https://educacao.uol.com.br/noticias/2017/08/28/mensalidade-de-curso-superior-no-brasil-custa-em-media-r-898-diz-estudo.htm .

[10] Considering SISU 2016.1 and 2016.2. In 2016.1 alone, the rate grows to 13.1 students per vacancy.

[11] Admissions processes to federal higher education institutions in Brazil are based, exclusively, on grades in one admission exam. Today, all federal institutions offer vacancies based on grades of the National Standardized Exam, mostly through a centralized admission system. Some of these institutions also offer part of their vacancies based on a specific exam elaborated by them. In any case, admissions are decided exclusively based on entrance exams, not taking into account high school performance.

universities come from the same demographic group (Brazilian Census 2010).

In order to improve equality in access to the federal tertiary education system, the government of Brazil approved law 12.711 on August 29^{th} 2012, the so-called the "Quota Law" (*QL*). It establishes that 50 percent of all vacancies in each major at each federal institution has to be reserved to students that attended all three years of secondary education in a public school. Moreover, there are subquotas, within these 50 percent, destined to racial and economic based minorities. Figure 2.2 shows an example of how *QL* was implemented in the State of Bahia. Take, for example, a major that offers 100 vacancies in the Federal University of Bahia. From these spots, 50 are reserved to students that attended all high school in a public school. Within these 50, 25 are reserved to public school students that belong to a family with per capita income of less than 1.5 minimum wage. Also, within these 50, 38 are reserved to black, mixed or indigenous students (non-white). The fraction of vacancies reserved to non-white students vary by state, according to the share of its population belonging to this demographic group.

By reserving 50 percent of vacancies at highly competitive institutions to public school students, *QL* increases incentives for public school attendance. Moreover, by establishing a national unique level of quotas for all federal institutions, it impacted differently each institution, depending on their pre-reform levels of quota adoption. For example, while some institutions already reserved as much as 50 percent of their vacancies to public school students (e.g. Federal University of Juiz de Fora - UFJF), others had no quotas at all (e.g. Federal University of Pernambuco - UFPE). In sum, the adoption of *QL* creates a quasi-experiment for testing how affirmative action in undergraduate education impacts high school students' outcomes and, by extension, their high schools and classmates.

2.3 Data and Variables

This paper uses data from multiple sources. First, it uses the Brazilian Census of Basic Education (CBS) from years 2008 to 2016. This is an administrative registry individual data of all students enrolled in primary and secondary schools in Brazil. It is collected yearly by the National Institute of Educational Studies and Research (INEP) of the Brazilian Ministry of Education and it is publicly available. The individual-module of the CBS contains basic demographic characteristics of students (e.g. gender, age, ethnicity) and unique individual and school identifiers. This allows for the construction of a panel dataset both at the individual and the school level across time. I select all individuals enrolled in the final year of primary education - 9^{th} grade - in year t from 2008 to 2015. Then, using the individual *id*, I link students from the CBS of cohort t with their own information in year $t+1$ from 2009 to 2016. This allows me to identify students that advanced to the first year of secondary education. Furthermore, I am able to identify whether individuals changed schools and, more importantly, whether they moved from their original educational system (from private to public or vice-versa).

Second, this paper uses detailed information on the adoption of AA quotas by all Brazilian public tertiary institutions. This dataset was collected by the author based on documents of the admission processes and on direct contact with the institutions. From this information, I constructed variable $Q_{u,2012}$, which measures the percentage of quotas at institution u of microregion m destined to students that attended all secondary education in a public school in year 2012, before the approval of the Quota Law. This allows me to understand how each institution in Brazil was exogenously exposed to QL, based on their pre-reform level of quota adoption. I, then, construct $Treat_{u,m} = 2(0.5 - Q_{u,2012})$. This means that, if the institution had no quotas in 2012 ($Q_{u,2012} = 0$), $Treat_{u,m} = 1$. On the other hand, if it already had 50 percent of reserved vacancies before the implementation of the law, then $Treat_{u,m} = 0$. If $0 < Q_{u,2012} < 0.5$, $Treat_{u,m}$ will assume a value between zero and one.

Although I directly observe how the QL reform affects all Brazilian 107 tertiary federal institutions, I do not know how students in the 9^{th} grade respond to changes in these different institutions. In order to obtain a proxy for that, I construct a measure of exposure to treatment based on the microregion m where the individual goes to school. First, I restrict my sample to 9^{th} graders that reside in a microregion where there is a federal *university*. Although there are 509 microregions in Brazil, there are only 50 with federal universities.[12] Table 2.1 presents descriptive statistics comparing the selected sample (hereafter *Sample 50*) with the whole population. *Sample 50* includes, generally, the larger and most developed microregions in Brazil. They encompass 45 percent of the total population, including all the State capitals and the Federal District (See Figure 2.3 for their location). The median student of *Sample 50* lives in a microregion where the median wage of the formal labor market is, on average, 25.5 percent higher than in the microregion of the median student of the whole population. Additionally, in *Sample 50*, 81.1 percent of 9^{th} graders attend a public school, compared to 86.5 percent in the population. Finally, note that, by living in a microregion with a federal university, students of *Sample 50* are more directly and clearly exposed to treatment, being more likely to be interested to attend college. In spite of that, all results from this paper remain robust for the complete population of 9^{th} graders in Brazil.

I construct two different measures of exposure to treatment for students of microregion m. These measures are based on size, quality and location of the institution.[13] The first,

[12]In these 50 microregions, there are 94 federal institutions, out of which 59 are universities. The other ones are smaller teaching centers and institutes.

[13]The rationale is that the larger the size and the quality index, more applications an institution receives,

and simpler measure, assumes that a student in microregion m is only exposed to treatment that occurred in institutions of the same microregion. According to the Brazilian Population Census of 2010, only 14 percent of college-enrollees moved, after age 14, to the municipality where they currently attend tertiary education. This suggests that students are mostly exposed to treatment through institutions in the microregion where they reside at age 14. Then, the first measure constructed is a weighted average of exposure to treatment of all federal institutions of region m, weighted by a measure of size and another of quality:

$$Treat_m^1 = \frac{\sum_{u=1}^{n} Quality_{u,m} x Size_{u,m} x Treat_{u,m}}{\sum_{u=1}^{n} Quality_{u,m} x Size_{u,m}}. \tag{2.1}$$

$Size_{u,m}$ is the number of new vacancies offered by the institution in 2012, while $Quality_{u,m}$ is a quality level index produced by INEP.[14] Note that if the microregion has only one federal university, $Treat_m^1 = Treat_{u,m}$.

In the second measure, I incorporate the possibility that individuals might be affected by institutions outside their microregion m. I use the Brazilian Population Census of 2010 and restrict data to individuals aged 18 to 25 that report to currently attend college. Then, I observe the municipality where individuals used to live *before college*.[15] Using this information, I construct the following measure:

$$Treat_m^2 = \theta_m x Treat_m^1 + \sum_{r \neq m} \theta_r Treat_r^1 \tag{2.2}$$

In this case, the variable $Treat_m^1$ for microregion m is weighted by the percentage of individuals that lived in m before college and keep doing so during college, θ_m. Additionally, it includes component $\sum_{r \neq m} \theta_r Treat_r^1$, where θ_r is the percentage of individuals that lived in m before college and moved to r (any of the other 49 microregions) for college. Take for example the case of the microregion of *Chapecó*, located in the State of *Santa Catarina*, south of Brazil. From all students that used to live in *Chapecó* before college and that attend a public tertiary institution, 81.2 percent do so in *Chapecó*. Additionally, 14.7 percent move to *Florianópolis*, the capital of *Santa Catarina*, 1.5 percent move to *Curitiba*, 1.5 percent to *Juiz de Fora* and 1 percent to *Pelotas*. Moves to all other cities are negligible (less than 0.5 percent). Therefore, variable $Treat^2_{Chapeco}$ is based on the weighted average that

and stronger is the general interest towards the institution.

[14] The index is called *Índice Geral de Cursos* and it varies from 1 to 5. Higher quality institutions receive more applications if compared to lower quality ones of the same size.

[15] Among all the students that attend college in the Census of 2010, 66 percent report to do so in the same municipality where they were born, while 20 report to have moved to their current city before age 14. The remaining 14 percent of students report to have moved from age 15 on wards. These are the ones I assume could have moved specifically for college.

attributes 81.2 percent for the microregion of *Chapecó* itself and 18.8 percent to all these other microregions, accordingly. Note that, out of the 50 microregions considered, θ_m for the median one is 96.8 percent. Therefore, while $Treat_m^2$ is conceptually a better measure, results are very similar to specifications that use $Treat_m^1$, instead. For the sake of simplicity, I only report results that consider $Treat_m^2$.

Figure 2.3 shows the distribution of variable $Treat_m^2$ and the location of microregions of *Sample 50*. First, note that the 50 microregions with a federal university are spread in different regions of Brazil. Second, note that they are relatively far from each other, minimizing concerns of spillover effects. Third, note that variable $Treat_m^2$ is sufficiently distributed between values zero and one. Although 20 out of 50 microregions are highly treated ($Treat_m^2 > 0.9$), 30 microregions have values of $Treat_m^2$ that vary from 0 to 0.9. It is this variation that allows for the causal identification of the effect of QL.

Finally, this paper uses microdata from the National Standardized Exam of High School (ENEM) from years 2011 to 2015. This dataset contains information on test scores of all students that took the exam in each year, their high school identifier and demographic characteristics. It is publicly available and produced by INEP. I use it to construct measures of school quality, as carefully explained in the Section 5 of this paper.

2.4 Strategic High School Choice

2.4.1 Empirical Strategy

I use an event-study design to study how the QL reform affects school choice. My main empirical model is:

$$Y_{imst} = \beta_t \sum_{t=2008}^{2015} Year_t Treat_m^j + \gamma X_{imst} + \delta X_{mt} + \alpha_{sm} + \alpha_t + \varepsilon_{imst}, j = 1,2 \tag{2.3}$$

where Y_{imst} is the outcome of student i, microregion m, school s and time t. The treatment variable $Treat_m^j$ defines how individual i was exposed to the QL reform, depending on the microregion where he or she attended 9^{th} grade, as explained in the previous section. This variable is interacted with a dummy for each year/cohort t from 2008 to 2015. The reform was announced in August of 2012. Therefore, the first cohort I expect to be influenced is exactly the one finishing 9^{th} grade in that year.[16] In this specification, years 2008 to 2010

[16]Note that to benefit from QL, students need to stay the full 3 years of high school in a public institution. Therefore, students already enrolled in secondary education cannot be affected

serve as pre-periods, 2011 is the baseline year and cohorts 2012 to 2015 are treated. The vector X_{imst} identifies individual controls (gender, age, ethnicity and urban status), while X_{mt} are time-varying microregion controls (size of formal sector, median wage of formal sector workers and share of college among formal sector workers). Finally, I include school fixed effects α_{sm} (which also absorb microregion fixed effects) and time-fixed effects α_t. Standard errors are clustered at the mesoregion level.[17]

The main identifying assumption for causal interpretation of parameters β_{2012}, β_{2013}, β_{2014} and β_{2015} is that dynamics in the outcome variable for treated and control units would have been similar in absence of the treatment. The presence of school-microregion fixed effects absorbs all unobservable time-invariant characteristics at school or microregion that might be correlated with the outcome. However, the existence of time-varying unobservable characteristics that are correlated with the outcome could still be a threat to causal identification. To minimize this concern, I include pre-periods 2008 to 2010. If pre-trends are parallel, I expect to find coefficients β_{2008}, β_{2009} and β_{2010} to be close to zero and insignificant. This would provide suggestive evidence that trends between treated and control microregions, in absence of treatment, would likely have been parallel also between 2011 and post-reform years.

Additionally, Table 2.2 presents results of regressions of different covariates on $Treat_m^j$ in 2011, the baseline year. Although not significantly different from zero, point estimates suggest the possible existence of difference in levels between treated and control microregions in 2011. Note that these differences are not necessarily a problem, as long as trends between groups are parallel.

2.4.2 Results

Tables 2.3 presents the baseline results from an OLS regression of equation 2.4. In Column (1), I control for time fixed effects only, while, in Column (2), I add school fixed effects. In column (3), I also control for individual characteristics and, in column (4), for microregion characterists. Note that my results are robust and remain very similar in all specifications. In my preferred model (Column 4), a full adoption of QL increases mobility of 9^{th} graders from a private to a public school by 2.5 p.p., in 2012, and by 4.5 p.p., in 2013. This represents an increase of 16 and 29 percent, respectively, in relation to the mean in baseline. Since the reform was approved on August 29^{th} 2012, the cohort of 9^{th} graders of 2012 did not have enough time to respond to changes, if compared to later cohorts. The academic year in Brazil goes from February to December and children need to enroll in public school by October of

[17] Brazil has, in total, 509 microregions (*Região Geográfica Imediata*) and 133 mesoregions (*Região Geográfica Intermediária*), defined by the Brazilian Institute for Geography and Statistics (IBGE). In my sample, I use 50 microregions, located in 45 different mesoregions.

the preceding year. Therefore, the 2012 cohort had approximately two months (September and October) to respond to the policy change. Later cohorts, on the other hand, had over a year to adjust. This is probably what explains the difference in magnitude between the estimates for year 2012 and later years. Finally, note the coefficients for years 2008, 2009 and 2010 are close to zero and insignificant, corroborating the identifying assumption of parallel trends between treated and control groups.

In Table 2.4, I study if the QL reform impacts mobility from public to private school. One could expect that, due to larger incentives for public school attendance, QL would decrease these movements in the transition between 9^{th} grade and secondary education. Results show that this is not the case. The reform has zero effect on these potential changes. This is possibly due to the fact that, before QL, these movements were already too low. In 2011, only 2.2 percent of 9^{th} graders that attended a public school moved to a private school in the following year. Results from Tables 2.3 and 2.4 are portrayed graphically in Figure 2.4.

Tables 2.5 and 2.6 show results analogous to Table 2.3 separately by ethnicity. As mentioned previously, in addition to the quotas destined to public school students, the QL reform has subquotas reserved to individuals self-declared as black, mixed or indigeneous (non-white). Therefore, this demographic group has higher incentives for strategic school choice. This pattern is confirmed in the results. According to Table 2.5 Column (4), a full adoption of QL increases moves from private to public school by 6.9 p.p. for non-whites in 2013, compared to 2.8 p.p. for whites and Asians-descendants. This is also shown graphically in Figure 2.5.

Finally, in Tables 2.7 and 2.8, I exploit heterogeneity of results with respect to students' school of origin. In Table 2.7, I split private schools in four quartiles of socioeconomic level, based on an index computed by INEP (*INSE 2011-2013*[18]). In 2013, a full adoption of QL increases movements of 9^{th} graders from private to public schools by 6.7 p.p. for the lowest-SES schools and by 0 p.p. for the highest-SES schools. Additionally, in Table 2.8, I divide schools into two categories based on a index of quality, also elaborated by INEP (*ENEM Escola*). This index is based on grades of high-school concluders in the National Standardized Exam of High School (ENEM) of 2011. Note that only a portion of middle schools are evaluated based on this index: the ones that also offer secondary education and the ones for which at least 50 percent of high-school concluders take the ENEM exam in 2011. The "Top Schools" in Table 2.8 are the ones with an index in the top half of the distribution. The "Other Schools" are the ones classified in the bottom half and the ones not classified, either because they do not offer high school or because less than 50 percent of their high-school concluders take the ENEM exam. The "Other with HS" refers to schools either

[18] INSE stands for *Índice de Nível Socioeconômico das Escolas de Educação Básica.*

classified in the bottom half or not classified, but that offer secondary education. Table 2.8 shows that a full adoption of QL increases movements of 9^{th} graders from private to public schools by 3.1 p.p. for the top schools and by 5.8 p.p. for the other (bottom) schools, in 2013. These results are shown, graphically, in Figures 2.6 and 2.7.

In sum, the adoption of QL increases significantly the strategic mobility of 9^{th} graders between private and public schools, in the transition between middle to secondary education. This is so because QL increases incentives for students to attend the full three years of high school at a public institution. The effect is stronger for non-whites, who benefit from additional subquotas if compared to whites, and for students coming from a low-SES and low-quality private school. Therefore, by increasing the *value* of public school, QL pushes some students from the private to the public system. These are individuals interested in public tertiary education, but that, most likely, benefit the least from private secondary school attendance. In contrast, students that stay in private school are the ones with a stronger preference for this type of institution, even after QL increased the value of public school for all individuals. These are, likely, individuals that benefit more from private-school attendance.

2.5 Effect on the High School System

Having established that QL increases mobility from private to public schools, some additional questions surface: does this exogenous inflow of students affect the quality of public high schools? If so, how? Are all schools equally affected? How do the effects spill over to students originally attending these schools? What about the quality of private schools that suffer from a substantial outflow of students? Finally, does QL impact the inequality between the public and the private high school systems? What about the inequality within each system? These are questions of extreme policy-relevancy.

As mentioned previously, Brazil is the 11^{th} most unequal country in the world and one key root of inequality is precisely its educational system. For example, in Figure 2.1, I showed how private schools perform considerably better than their public counterparts. Moreover, inequality in access to high-quality public tertiary education is substantial. The QL reform aims specifically at reducing inequality at this later stage of one's educational trajectory. Most critics of AA in higher education would argue that such policies have limited effects, since inequality in college access has its roots much earlier in the student's academic career. In this context, it is extremely important to assess whether QL impacted the high school system as a whole and, importantly, whether it had any distributional effects.

In this version of the manuscript, I focus on assessing whether QL affects the quality of public schools that receive these previously private-school kids.[19] To do so, I construct a panel dataset at school level, including only public high schools. Then, I compute variable $SharePriv_{pmy}$, which is the share of incoming students in the first year of secondary education at public school p, microregion m and year y, that came from a 9^{th} grade at a private school. Due to data constraints for the quality index, I focus on cohorts of first year high school students of years 2010 and 2011, for pre-reform periods, and of 2013 and 2014, for post-reform. Note that, in comparison to the previous section, there is a slight change of notation. For example, 9^{th} graders of time *t=2009* are now first year high school students of year *y=2010*.

Table 2.9 presents the distribution of variable $SharePriv_{py}$ for the cohort of 1^{st} year high school students of 2010. Its median share is 1.3 percent, the 90^{th} percentile is 8.2 and the 99^{th} percentile is 42 percent. This means that, from the 6443 of Brazilian public high schools of *Sample 50*, the vast majority receives none or only a very small number of students that went to a private middle school. I, then, select only the top quartile and top decile (hereafter *Top Quart* and *Top Dec*) of the sample based on variable $SharePriv_{pm,2010}$. These are the public schools that, in pre-periods, received most of the 9^{th} graders from a private school. I, then, estimate the following model, for each of these sample of schools:

$$Y_{pym} = \beta_y \sum_{y=2010,2011,2013,2014} Year_y Treat^j_m + \alpha_{pm} + \alpha_y + \varepsilon_{mpy}, j = 1, 2, \tag{2.4}$$

where Y_{pym} is the outcome for public school p, microregion m and year y. The treatment variable, $Treat^j_m$, at microregion level, is interacted with years 2010, a pre-period, 2013 and 2014, post-reform periods. Year 2011 is the baseline in this case. The terms α_{pm} controls for school fixed effects (while also absorbing microregion fixed effects) and α_y controls for year fixed effects. Standard errors ε_{mpy} are clustered at mesoregion level and observations are weighted by the number of students in each school.

Table 2.10 shows results for the outcome $SharePriv_{pmy}$. Considering all public schools, the adoption of QL increases the share of previously private school students by 1.4 percent in 2013 and by 2.0 p.p. in 2014. In contrast, focusing on the sample *Top Dec*, the effect of QL is 4.7 p.p. in 2013 and 7.5 p.p. in 2014. I, then, investigate the effect of QL on school quality. To measure quality, I rely on grades of high school concluders at the National Standardized Exam of High School (ENEM). For each cohort of first year high school students in year y, I

[19] I am currently working on the additional questions I posed in the first paragraph. Specifically, in the next version of this manuscript, I plan to investigate whether the effect QL spills over to students originally attending public schools. Also, I plan to assess how it affects the quality of private schools. This will allow me to have a more comprehensive understanding of how QL affects the high school system as a whole.

obtain their average grades in the ENEM exam of year *y+2*, when they conclude high school. Results from Table 2.11 show that, for all schools, the effect on quality is zero in 2013 and around 0.04 standard deviations in 2014, although the point estimate is not significant. For sample *Top Dec*, alternatively, the effect of *QL* is 0.03 standard deviation in 2013 and 0.08 in 2014. Most importantly, Table 2.11 presents convincing evidence that *QL* has a positive effect on the quality of public schools, especially for the ones in the *Top Dec* sample, which received considerably more private school students. Taken together, these results suggest that the mobility of private school students to public schools induced by *QL* is an important driver of this increase in quality.[20]

However, it is important to point out that this effect is heterogeneous along the distribution of schools. Note that schools in each sample are very different in terms of pre-levels of quality. In the complete sample, the average quality of public schools in 2010 is -0.26 (as measured in a standardized normal with mean zero and standard deviation 1). In the *Top Dec* sample, the average is 0.13. This means that, when students decide to move from a private to a public school, they tend to select originally better public schools than the average one. Moreover, the *QL* reform seems to reinforce this pattern. It mostly increases mobility to schools in the *Top Dec* sample. Additionally, it increases quality disproportionately at these better public schools. Although additional research is needed, these results suggest that *QL* increases overall quality of public schools. Yet, it seems to also increase inequality across schools within the public system. In future versions of this manuscript, I plan to dig more into this question and to understand whether this positive effect on quality spills over to students originally attending the public schools.

2.6 Concluding Remarks

In this paper, I studied how the adoption of a national affirmative action initiative in the Brazilian federal higher education system impacted students' outcomes before college. In 2012, the government approved the national Quotal Law (*QL*), which reserved 50 percent of all vacancies in each major and institution to students that attended all three years of high school in a public institution. By changing incentives for admissions to the most competitive tertiary institutions in the country, the government encouraged (unintendedly or not) public high school attendance.

I find that a full adoption of *QL* increases strategic mobility from private to public schools by 29 percent. Moreover, I find that this effect is stronger for non-whites, who benefit from an additional subquota, and for students originally attending low-SES and low-quality private

[20] In the next version of this manuscript, I plan to test this directly through an IV-approach.

schools. Moreover, I investigate how this exogenous influx of private school students affect the quality of public schools. First, I find the these students choose to move to public schools that are originally better than the average one. Second, I find that, even if they come from low-SES and low-quality private schools, QL has a positive effect on public school quality. However, this effect is concentrated on a small number of schools that were already better than the average one. Therefore, it seems that the mobility from private to public schools induced by QL positively affects the quality, but simultaneously increases inequality within the public school system.

There are still many questions left unanswered, which I plan to address in future versions of this manuscript. Is this positive effect on quality just a result of a composition effect or does it also spill over to students originally attending public schools? What dimensions do these possible spillover effects influence? For example, in addition to grades, I plan to look at high school persistence, graduation and college application decisions. How are private schools affected? Finally, what are the overall consequences of this policy for the high school system? Preliminary results for the public school system suggest a positive effect on integration for a small number of public schools, but an increase in segregation between these schools and the rest of the public system. Additional research is needed to clarify such points and for a better understanding of the overall effects of QL and, more generally, of whether AA policies in higher education are capable to reduce inequality in the basic education system.

2.7 Figures and Tables

Figure 2.1: Private vs Public Schools

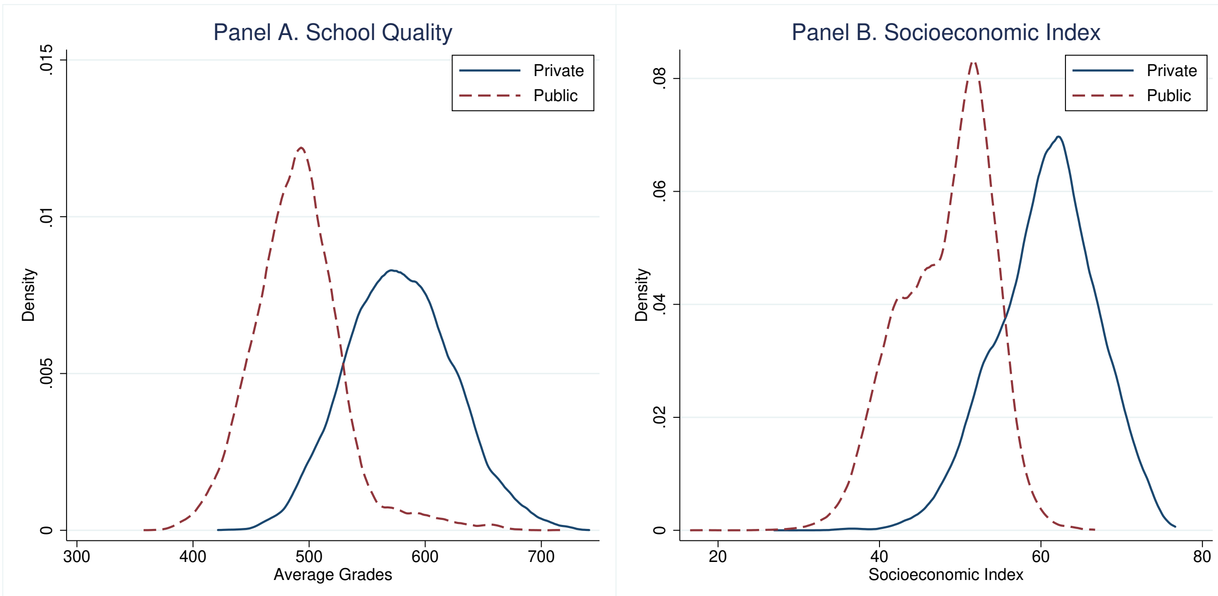

Notes: This figure shows the distribution of private and public schools in Brazil, according to a quality and a socioeconomic index. It includes all schools that offer the second cycle of primary (grades 5-9) and secondary education, according to the Census of Basic Education of 2011. The measure of school quality comes from average grades in the National Standardized Exam of High School (ENEM Escola 2011). The Socioeconomic Index (INSE 2011-2013) is computed by INEP in a continuous scale with mean 50 and standard deviation 10.

Figure 2.2: Example of the Quota Law for the State of Bahia

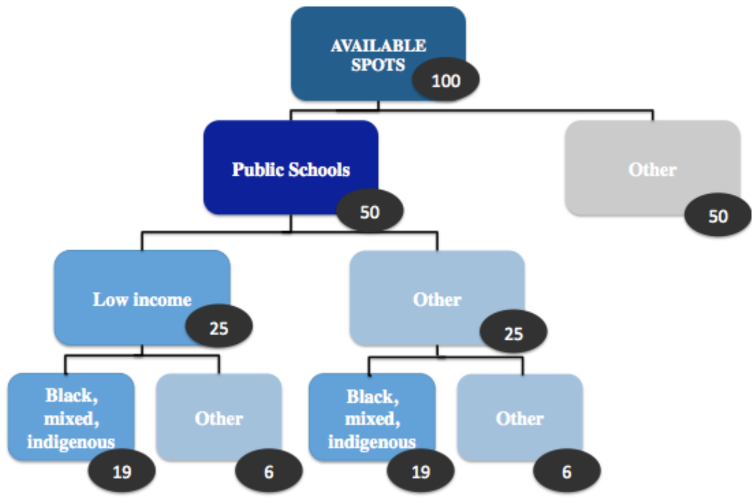

Notes: From every 100 vacancies offered by each major in a federal institution, 50 are reserved to students that attended all three years of secondary education in a public school, while the other 50 are open for regular competitive entrance. From the vacancies reserved to public school students, half are also reserved to individuals from families of per capita income of less than 1.5 minimum wage. Finally, the State of Bahia has 77 percent of non-white individuals according to the Census of 2010. Therefore, 77 percent of the reserved vacancies has to be destined to the non-white. **Source:** own elaboration based on Law 12.711/2012.

Figure 2.3: Distribution of Treatment Variable

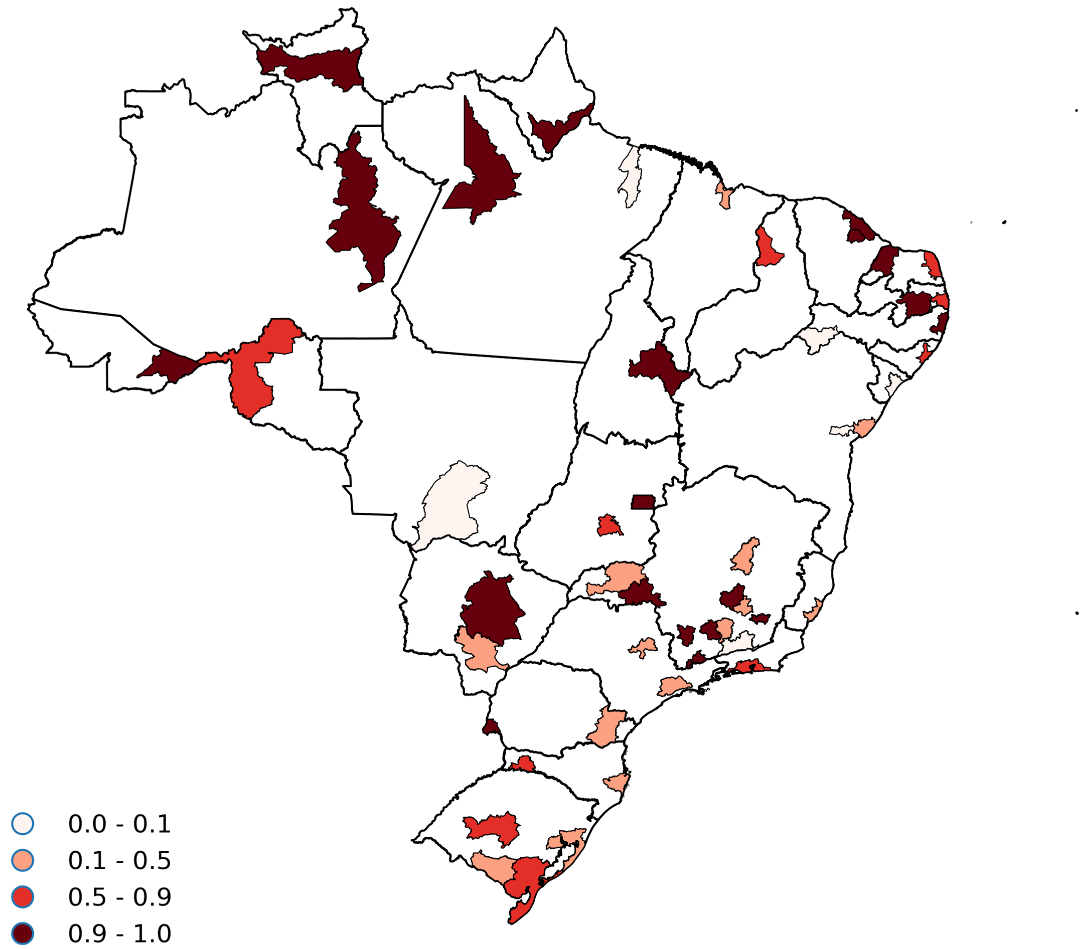

Notes: This map shows the location of the 50 microregions in Brazil with a federal university, which I referred to as *Sample 50* in the main text. The microregions are coloured according to their level of treatment, which can vary from zero (not-treated) to 1 (fully-treated).

Figure 2.4: Estimate of Treatment Effects between School Systems

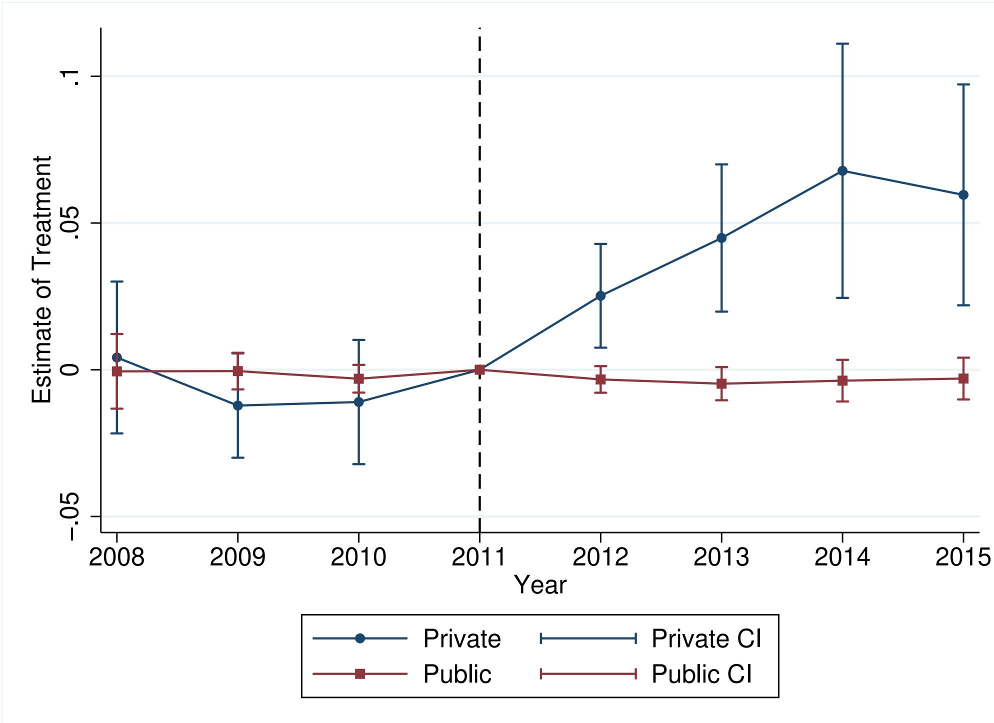

Notes: In this Figure, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders to secondary school. *Private* refers to movements from private to public schools and *Public* to movements from public to private institutions. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. CI are 95% confidence intervals. Specifications include year and school FE, individual and microregion time-varying controls.

Figure 2.5: Estimate of Movements from Private to Public Schools by Ethnicity

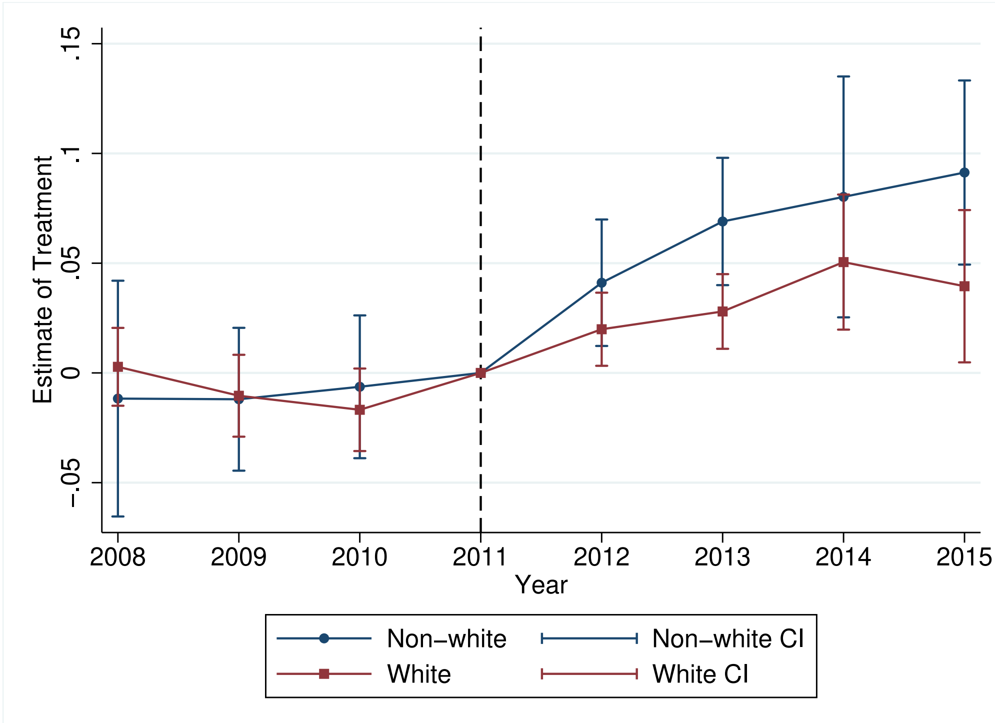

Notes: In this Figure, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school. *Non-white* refers to individuals self-declared as black, indigenous or mixed and Whites to individuals self-declared as whites or Asians. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. CI are 95% confidence intervals. Specifications include year and school FE, individual and microregion time-varying controls.

Figure 2.6: Estimate of Movements from Private to Public Schools by School of Origin Socioeconomic Status

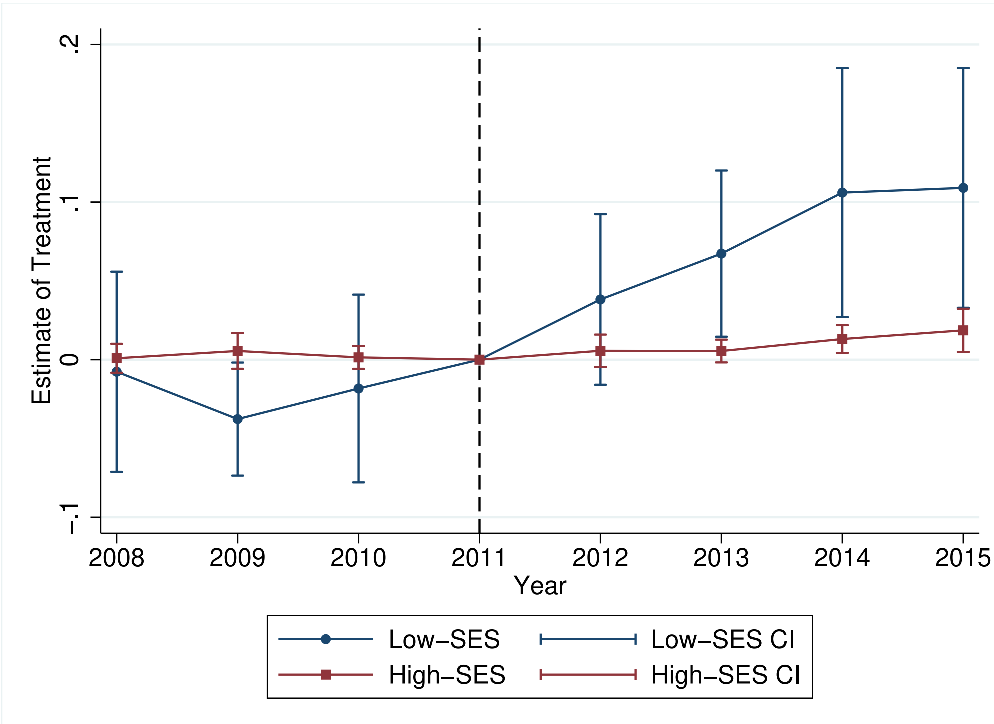

Notes: In this Figure, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school. *Low-SES* refers to origin schools in the bottom quartile of the distribution of socioeconomic status according to an index computed by the National Institute of Educational Studies and Research in Brazil (so-called INSE 2011-2013). *High-SES* refers to origin schools in the top quartile. CI are 95% confidence intervals. Specifications include year and school FE, individual and microregion time-varying controls.

Figure 2.7: Estimate of Movements from Private to Public Schools by School of Origin Quality Index

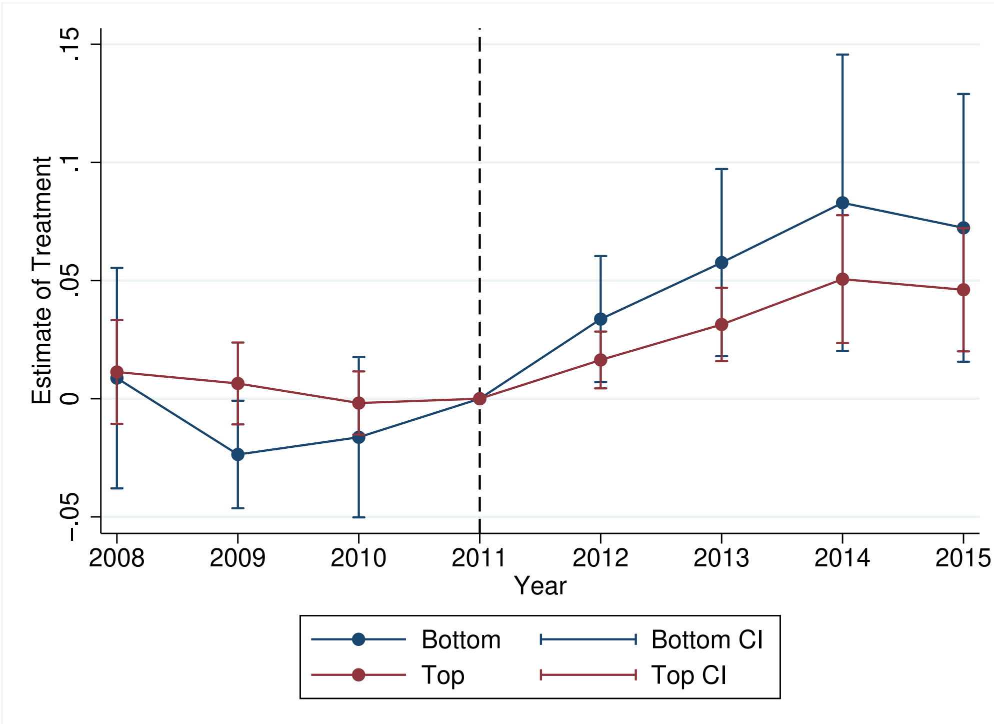

Notes: In this Figure, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school. The Quality Index was based on grades at the National Standardized High-School Exam of 2011 and computed by INEP (*Enem Escola 2011*). *Top* refers to private schools classified in the top half of the distribution of test scores and *Bottom* refers to all other schools, both classified in the bottom-half or not classified at all. CI are 95% confidence intervals. Specifications include year and school FE, individual and microregion time-varying controls.

Table 2.1: Descriptive Statistics

Cohort	Observations		Median Wages		Share College	
	All	*Sample 50*	*All*	*Sample 50*	*All*	*Sample 50*
2008	2974583	1339030	11060	13405	30.4	32.8
2009	3042326	1373501	11732	13776	30.4	32.8
2010	2998554	1352048	12477	14720	30.3	32.8
2011	2972360	1346985	13273	15864	30.7	32.9
2012	2973644	1349088	14484	16625	31.5	35.1
2013	2990971	1360930	15410	17581	31.5	35.7
2014	2856323	1310270	15960	18285	31.5	33.7
2015	2722276	1224427	15471	17583	31.7	33.9
Total	23531037	10656279	13561	17015	31.0	33.7
Cohort	**Public School**		**Private to Public**		**Public to Private**	
	All	*Sample 50*	*All*	*Sample 50*	*All*	*Sample 50*
2008	88.1	83.5	16.7	15.9	1.8	1.8
2009	87.8	83.0	16.2	15.3	2.0	2.0
2010	87.1	81.9	16.0	15.2	2.0	2.1
2011	86.4	81.0	16.4	15.4	2.0	2.2
2012	86.1	80.4	18.5	17.6	1.8	1.8
2013	85.8	80.2	19.3	18.4	1.8	1.9
2014	85.1	79.3	21.0	20.3	1.6	1.6
2015	85.3	79.4	21.9	20.9	1.5	1.5
Total	86.5	81.1	18.3	17.5	1.8	1.9

Notes: This table compares descriptive statistics for the Sample of 50 microregions I use for the baseline analysis, versus the complete population. *Median Wages* and *Share College* are, respectively, the median wages and proportion of individuals with a college degree among the formal workers in the microregion of the median 9^{th} grade student. *Public School* is the proportion of 9^{th} graders attending a public school, while *Private to Public* and *Public to Private* refer to moves from one system to other, between 9^{th} grade of primary to the 1^{st} year of secondary education.

Table 2.2: Regression on Treatment Status in 2011

	$Treat_m$	Standard Errors	Observations
Public School	-0.0436	*(0.0400)*	1346985
Private to Public	0.0366	*(0.0300)*	221604
Public to Private	0.000154	*(0.0105)*	811737
Log Median Wage	-0.403	*(0.291)*	1346985
Log N Formal	-0.350	*(0.455)*	1346985
Share College	-0.135	*(0.0881)*	1346985

Notes: This contains results of regressions of covariates on the variable for treatment $Treat_m$ in year 2011, the baseline. *Public School* is the proportion of 9^{th} graders attending a public school, while *Private to Public* and *Public to Private* refer to moves from one system to other, between 9^{th} grade of primary to the 1^{st} year of secondary education. *Log Median Wages*, *Log N Formal*, *Share College* are, respectively, the log of median wages, the log of the number of workers and the proportion of individuals with a college degree, among the formal workers in the microregion.

Table 2.3: Moves from Private to Public School

	(1)	(2)	(3)	(4)
$Treat_m$.2008	-0.00215	-0.00177	-0.00140	0.00418
	(0.0109)	(0.0132)	(0.0132)	(0.0132)
$Treat_m$.2009	-0.0132	-0.0165	-0.0162	-0.0122
	(0.0113)	(0.0110)	(0.0108)	(0.00907)
$Treat_m$.2010	-0.0177	-0.0126	-0.0123	-0.0110
	(0.0120)	(0.0100)	(0.00998)	(0.0108)
$Treat_m$.2012	0.0308***	0.0292***	0.0290***	0.0252***
	(0.0110)	(0.0101)	(0.0100)	(0.00902)
$Treat_m$.2013	0.0603***	0.0505***	0.0502***	0.0449***
	(0.0205)	(0.0156)	(0.0155)	(0.0128)
$Treat_m$.2014	0.0801**	0.0752***	0.0749***	0.0678***
	(0.0344)	(0.0264)	(0.0261)	(0.0221)
$Treat_m$.2015	0.0682**	0.0702***	0.0696***	0.0596***
	(0.0292)	(0.0222)	(0.0219)	(0.0192)
N	1724125	1724080	1724080	1724080
Adj. R2	0.027	0.232	0.233	0.233
Time FE	Yes	Yes	Yes	Yes
School FE		Yes	Yes	Yes
Individual Controls			Yes	Yes
Region Controls				Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.4: Moves from Public to Private School

	(1)	(2)	(3)	(4)
$Treat_m$.2008	-0.000689	-0.000865	-0.00101	-0.000546
	(0.00607)	(0.00620)	(0.00621)	(0.00649)
$Treat_m$.2009	0.0000132	-0.000551	-0.000722	-0.000456
	(0.00279)	(0.00292)	(0.00294)	(0.00318)
$Treat_m$.2010	-0.00280	-0.00304	-0.00317	-0.00307
	(0.00237)	(0.00253)	(0.00254)	(0.00241)
$Treat_m$.2012	-0.00260	-0.00295	-0.00297	-0.00330
	(0.00261)	(0.00244)	(0.00241)	(0.00231)
$Treat_m$.2013	-0.00466	-0.00513*	-0.00481*	-0.00475
	(0.00298)	(0.00279)	(0.00267)	(0.00288)
$Treat_m$.2014	-0.00371	-0.00392	-0.00363	-0.00372
	(0.00424)	(0.00394)	(0.00378)	(0.00362)
$Treat_m$.2015	-0.00292	-0.00325	-0.00307	-0.00300
	(0.00419)	(0.00404)	(0.00383)	(0.00363)
N	6348207	6348152	6348152	6348152
Adj. R2	0.008	0.050	0.052	0.052
Time FE	Yes	Yes	Yes	Yes
School FE		Yes	Yes	Yes
Individual Controls			Yes	Yes
Region Controls				Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a public school to a private secondary school. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.5: Moves from Private to Public School - Non-white Students

	(1)	(2)	(3)	(4)
$Treat_m.2008$	-0.0410	-0.00934	-0.00962	-0.0117
	(0.0271)	(0.0257)	(0.0258)	(0.0274)
$Treat_m.2009$	-0.0169	-0.00919	-0.00928	-0.0120
	(0.0192)	(0.0163)	(0.0164)	(0.0166)
$Treat_m.2010$	-0.0223	-0.00705	-0.00695	-0.00632
	(0.0180)	(0.0156)	(0.0157)	(0.0166)
$Treat_m.2012$	0.0548***	0.0481***	0.0480***	0.0411***
	(0.0184)	(0.0132)	(0.0133)	(0.0147)
$Treat_m.2013$	0.0936***	0.0730***	0.0731***	0.0690***
	(0.0312)	(0.0212)	(0.0213)	(0.0148)
$Treat_m.2014$	0.0952**	0.0853**	0.0853**	0.0802***
	(0.0453)	(0.0337)	(0.0338)	(0.0280)
$Treat_m.2015$	0.0992**	0.0977***	0.0977***	0.0913***
	(0.0369)	(0.0263)	(0.0264)	(0.0214)
N	269836	269396	269396	269396
Adj. R2	0.047	0.255	0.255	0.255
Time FE	Yes	Yes	Yes	Yes
School FE		Yes	Yes	Yes
Individual Controls			Yes	Yes
Region Controls				Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school, for students self-declared as black, mixed and indigenous only (non-white). Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.6: Moves from Private to Public School - White and Asian Students

	(1)	(2)	(3)	(4)
$Treat_m$.2008	-0.00733	-0.00241	-0.00230	0.00278
	(0.0139)	(0.00984)	(0.00982)	(0.00905)
$Treat_m$.2009	-0.0143	-0.0150	-0.0150	-0.0104
	(0.0113)	(0.0107)	(0.0107)	(0.00952)
$Treat_m$.2010	-0.0248*	-0.0176*	-0.0176*	-0.0168*
	(0.0124)	(0.00979)	(0.00979)	(0.00959)
$Treat_m$.2012	0.0262**	0.0215**	0.0214**	0.0199**
	(0.0121)	(0.00953)	(0.00953)	(0.00850)
$Treat_m$.2013	0.0335**	0.0333***	0.0330***	0.0280***
	(0.0154)	(0.00942)	(0.00944)	(0.00868)
$Treat_m$.2014	0.0530**	0.0571***	0.0569***	0.0505***
	(0.0211)	(0.0177)	(0.0177)	(0.0157)
$Treat_m$.2015	0.0398**	0.0494***	0.0491***	0.0395**
	(0.0191)	(0.0177)	(0.0177)	(0.0177)
N	552670	552414	552414	552414
Adj. R2	0.019	0.201	0.202	0.202
Time FE	Yes	Yes	Yes	Yes
School FE		Yes	Yes	Yes
Individual Controls			Yes	Yes
Region Controls				Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school, for students self-declared as white or Asian. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.7: Moves from Private to Public School - By School Socioeconomic Level

	Quart 1	Quart 2	Quart 3	Quart 4
$Treat_m.2008$	-0.00766	-0.0146	-0.00537	0.000897
	(0.0324)	(0.0159)	(0.0170)	(0.00469)
$Treat_m.2009$	-0.0377**	-0.0213	-0.0168	0.00552
	(0.0183)	(0.0181)	(0.0112)	(0.00577)
$Treat_m.2010$	-0.0183	-0.0130	-0.0181**	0.00146
	(0.0304)	(0.0154)	(0.00785)	(0.00371)
$Treat_m.2012$	0.0382	0.0478***	0.00762	0.00564
	(0.0276)	(0.0162)	(0.00946)	(0.00524)
$Treat_m.2013$	0.0673**	0.0539***	0.0326***	0.00549
	(0.0269)	(0.0152)	(0.00909)	(0.00369)
$Treat_m.2014$	0.106**	0.0807***	0.0613***	0.0131***
	(0.0403)	(0.0207)	(0.0148)	(0.00448)
$Treat_m.2015$	0.109***	0.0772**	0.0495**	0.0186**
	(0.0388)	(0.0299)	(0.0220)	(0.00700)
N	220621	285117	329974	504639
Adj. R2	0.144	0.105	0.097	0.068
Time FE	Yes	Yes	Yes	Yes
School FE	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Region Controls	Yes	Yes	Yes	Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school, by school socioeconomic status. *Quart 1* refers to origin schools in the bottom quartile of the distribution of socioeconomic status according to an index computed by INEP (so-called INSE 2011-2013). *Quart 4* refers to origin schools in the top quartile. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.8: Moves from Private to Public School - By School Quality

	Top Schools	Other Schools	Other with HS
$Treat_m$.2008	0.0113	0.00871	0.00865
	(0.0112)	(0.0238)	(0.0214)
$Treat_m$.2009	0.00645	-0.0236**	-0.0301***
	(0.00884)	(0.0116)	(0.0102)
$Treat_m$.2010	-0.00183	-0.0163	-0.0146
	(0.00683)	(0.0173)	(0.0176)
$Treat_m$.2012	0.0164**	0.0337**	0.0382***
	(0.00613)	(0.0136)	(0.0136)
$Treat_m$.2013	0.0314***	0.0576***	0.0502**
	(0.00792)	(0.0202)	(0.0209)
$Treat_m$.2014	0.0506***	0.0829**	0.0906***
	(0.0138)	(0.0320)	(0.0324)
$Treat_m$.2015	0.0461***	0.0723**	0.0886**
	(0.0133)	(0.0289)	(0.0359)
N	812115	911964	474467
Adj. R2	0.082	0.238	0.117
Time FE	Yes	Yes	Yes
School FE	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes
Region Controls	Yes	Yes	Yes

Notes: In this Table, I plot the estimates of the treatment effects of QL on movements of 9^{th} graders of a private school to a public secondary school, by school quality index. The Quality Index was based on grades at the National Standardized High-School Exam of 2011 and computed by INEP (*Enem Escola 2011*). *Top* refers to private schools classified in the top half of the distribution of test scores and *Other Schools* refers to all other schools, both classified in the bottom-half or not classified at all. *Other with HS* refer to schools either in the bottom half or not classified at all, but that offer secondary education. Year 2011 is the baseline, 2008 to 2010 are pre-periods and 2012 to 2015 are treated periods. Standard errors are shown in parenthesis and are clustered at the mesoregion level. Individual controls include gender, age, ethnicity and urban status. Microregion controls include log of number of workers in the formal sector, log of yearly wage of the median worker in the formal sector and share of individuals with college among formal workers in the region.

Table 2.9: Distribution of $SharePriv_{pm,2010}$

	All	Top Quart	Top Dec
1%	0.00	2.86	7.16
5%	0.00	3.06	7.34
10%	0.00	3.30	7.54
25%	0.00	4.00	8.62
50%	1.30	5.98	12.41
75%	4.01	10.45	19.30
90%	8.62	19.30	37.50
95%	15.27	33.06	48.62
99%	42.17	64.35	79.07
Maximum	94.51	9.82	94.51
Mean	3.76	11.20	17.76
Std. Dev.	7.81	94.51	14.38
Obs	6443	1612	645

Notes: $SharePriv_{pm,2010}$ is the share of students in the first year of high school at school p that came from 9^{th} grade at a private institution. *All* contains all 6443 public schools in the sample; *Top Quart* only schools in the top quartile of $SharePriv_{p,2010}$ and *Top Dec* only the ones in the top decile. Observations are weighted by number of students.

Table 2.10: Effect on $SharePriv_{pmy}$

	All	Top Quart	Top Dec
$Treat_m$.2010	-0.00175	-0.00318	0.00444
	(0.00237)	(0.00428)	(0.00613)
$Treat_m$.2013	0.0143**	0.0273***	0.0474***
	(0.00581)	(0.00890)	(0.0143)
$Treat_m$.2014	0.0200**	0.0393***	0.0747***
	(0.00822)	(0.0135)	(0.0214)
N	26045	6323	2519
Adj. R2	0.881	0.865	0.850
Average Baseline	0.038	0.112	0.178
Time FE	Yes	Yes	Yes
School FE	Yes	Yes	Yes

Notes: $SharePriv_{pmy}$ is the share of students in the first year of high school at school p that came from 9^{th} grade at a private institution. *All* contains all 6443 public schools in the sample; *Top Quart* only schools in the top quartile of $SharePriv_{pm,2010}$ and *Top Dec* only the ones in the top decile. Standard errors in parenthesis are clustered at mesoregion level and observations are weighted by number of students.

Table 2.11: Effect on School Quality

	All Schools			Schools with Participation >0.5		
	All	*Top Quart*	*Top Dec*	*All*	*Top Quart*	*Top Dec*
$Treat_m.2010$	-0.0117	-0.0237	-0.0205	-0.00735	-0.00833	-0.0265
	(0.0121)	(0.0178)	(0.0260)	(0.0169)	(0.0205)	(0.0300)
$Treat_m.2013$	-0.00165	-0.0117	0.0298	-0.00586	-0.00788	0.0308
	(0.0192)	(0.0212)	(0.0281)	(0.0204)	(0.0238)	(0.0282)
$Treat_m.2014$	0.0359	0.0400	0.0779**	0.0225	0.0393	0.0744**
	(0.0283)	(0.0281)	(0.0351)	(0.0275)	(0.0290)	(0.0364)
N	25479	6271	2499	12706	4563	2097
Adj. R2	0.900	0.951	0.964	0.949	0.962	0.968
Average Baseline	-0.26	-0.06	0.13	-0.16	0.00	0.20
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
School FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: School Quality is measured by the average grade of the cohort in the National Standardized Exam of High School (ENEM). Cohort of first year high school students of 2013, for example, take the EXAM in 2015, near completion of secondary school. *All Schools* involve all schools with at least one student found in the ENEM. *Schools with Participation >0.5* involve only schools in which at least 50 percent of students take the ENEM. *All* contains all 6443 public schools in the sample; *Top Quart* only schools in the top quartile of $SharePriv_{pm,2010}$ and *Top Dec* only the ones in the top decile. Standard errors in parenthesis are clustered at mesoregion level and observations are weighted by number of students.

Chapter 3

Correction Methods for Intergenerational Mobility Estimates

3.1 Introduction

The estimation of the intergenerational elasticity (IGE) ideally requires complete income series for two generations of fathers and offspring for the construction of lifetime incomes. However, it is rare to find this ideal setting in empirical applications. What we observe, instead, are snapshots of data that contain a limited number of income observations in the life cycle. If those income snapshots do not mimic lifetime incomes, we have the so-called lifecycle bias. The literature has proposed different methods to address this issue, which can be classified into two alternative approaches. One option is to formulate an errors-in-variables model that formalizes the relation between between (observed) annual and (typically unobserved) lifetime income. For example, the generalized errors-in-variables model proposed by Haider and Solon (2006) is based on the intuition that for two individuals with different incomes trajectories, there will nevertheless exist an age at which the difference between their annual incomes equals the difference between their lifetime incomes. An alternative option is to directly model the income process that determines the relation between annual and lifetime incomes, and to predict lifecycle profiles and lifetime incomes from partially observed profiles. For example, Vogel (2007), Hertz (2007) and Creedy (1988) predict income growth based on observable characteristics, such as education or occupation.

In this paper, we use uniquely long series of Swedish and simulated data from a calibrated income process for the U.S. (Guvenen, 2009) to provide evidence on whether and why these correction methods reduce lifecycle bias, and to propose a more comprehensive way to estimate the IGE with incomplete income data. First, we analyze the key components of the income process that affect intergenerational estimators: (i) income growth explained by observable characteristics; (ii) transitory noise and (iii) unexplained income growth. In Section 3.3, we document that there is strong heterogeneity in income growth among individuals with similar observable characteristics, and that this heterogeneity relates systematically – and can be predicted by – parental background. Within each educational or occupational group, sons with fathers in the top income quartile have substantially steeper income growth, but lower initial incomes. These differences are sizable. For example, managers from high-income fathers tend to have lower income in their late 20s, but nearly 30 percent higher incomes around age 40 than managers from less well-off families. These findings contribute to a wider debate on the properties of income processes, supporting the argument that (residual) income grows at an individual-specific and deterministic rate. Because this individual heterogeneity relates systematically to parental income, it also represents an important obstacle for empirical work in the intergenerational literature.

We then analyze the most popular correction methods in Section 3.4, studying how they address these key components of the income process. Most intergenerational studies base their empirical strategy on the generalized errors-in-variables model (GEIV) proposed by Haider and Solon (2006), whose slope coefficient captures how differences in log annual income map into differences in log lifetime income. While properties of the income process are not explicitly considered, they inform the model's errors-in-variables assumptions. Its key prediction is that lifecycle bias can be minimized by measuring incomes at some age around midlife. However, the exact age at which bias is minimized is not known in applications. Table A.1 lists a number of recent papers, and how measurement error has been addressed in each of them. Most applications relate to the GEIV model in an informal way, as motivation for the use of income averages at some age. However, the chosen age range is limited by data availability and varies considerably across studies. For example, Chetty et al. (2014) use a 2-year average of offspring income around ages 29 to 32 (finding an IGE of 0.34), while Mazumder (2016) uses an 11-year average centered at age 40 (finding an IGE of 0.66). Such age variation is problematic as intergenerational elasticity estimates can vary substantially with the exact age at measurement.[1] An errors-in-variables model can therefore minimize lifecycle bias, but large biases are likely to remain.

An alternative approach is to estimate the income process itself. In this approach, researchers do not simply construct income averages from the observed snapshots, but use them to deduct the likely shape of an individual's profile in the unobserved age range. Vogel (2007), Hertz (2007) and Creedy (1988) each proposed methods that explicitly model key components of the income process. Vogel (2007) and Hertz (2007) allow for unobserved

[1]Consistent with the standard implication from the Mincer equation that age-earnings profiles diverge over age (see Heckman, Lochner and Todd, 2006)

individual heterogeneity in the level of income (i.e. individual fixed effects) and observed individual heterogeneity in income growth as explained by personal characteristics (such as education or occupation). However, we find that even conditional on those characteristics, children from high-income parents tend to have steeper income growth. As a consequence, these methods are still sensitive to the age at which income is being measured, and cannot eliminate lifecycle bias. The problem is likely to extend to the estimation of mobility trends, an issue that we return to below. The correction by Creedy (1988) does account for both observable and unobservable determinants of individual-specific income growth, but does not explicitly account for the influence of transitory income shocks. In sum, the three approaches capture important aspects of the income process, but neither captures the three key properties as illustrated in Section 3.3.

Based on these insights, we propose a new "*lifecycle estimator*" for estimating the IGE in incomplete income data, which exploits the available income information more fully. We aim to combine the strengths of previously proposed methods, while addressing its key limitations. In a first step, we estimate income profiles based on observable characteristics such as age and education, and individual fixed effects (similarly as Vogel 2007 and Hertz 2007). However, we introduce additional sources of heterogeneity, to account for the fact that children from high-income families tend to have steeper slopes even conditional on those observables. We avoid estimating an individual-specific growth rate for each individual, which would result in noisy estimates of the overall profile and lifetime incomes. Instead, our preferred specification allows for individual slopes to vary with parental characteristics or – if these are unobserved – with income levels. We can thereby capture one of the key insights from Creedy (1988), that income profiles tend to "fan out" over age, so that individuals with higher income *levels* also tend to have higher income *growth* over their life cycle.

We show that this lifecycle estimator yields substantially better results, with estimates of the intergenerational elasticity fluctuating closely around the benchmark estimate based on long-run incomes. Moreover, the estimator performs well even when there are only very few (as low as two) income observations available for each person. The method promises, therefore, to be applicable in a wide range of data settings. In contrast to current practice, the method exploits all the available income information in the data, and can be used even when incomes around midlife are not directly observed for the child generation.[2]

In addition to this lifecycle estimator, we present a "standardized" errors-in-variables model. It is similar to the generalized errors-in-variables model proposed by Haider and Solon

[2]We plan to adapt this method also for the estimation of lifetime incomes in the father generation. Its application to the parent generation is not straightforward because classical measurement error does not affect intergenerational estimates if occurring on the left-hand side (the child generation), but it will cause attenuation bias if occurring in the right-hand side variable (the parent generation).

(2006), but based on standard moments of the income process that might be more readily estimable in applications. Compared to the classical error-in-variables model, it accounts for the signal-to-noise ratio but also the fanning out of income profiles over age. Our objective underlying either of the two methods is to move beyond the simple rule-of-thumb approach in the current literature, and to enable more targeted corrections of intergenerational mobility estimates.

Our paper relates to an extensive and growing literature on intergenerational mobility that has attempted to measure the levels and trends of the IGE in different countries (see Solon (1999), Black and Devereux (2011) and Jäntti and Jenkins (2015) for reviews). Specifically, it contributes to the literature on measurement error in intergenerational mobility estimates. The earlier papers in this literature focused on classical measurement error from incomplete income data for fathers (Atkinson 1980, Solon 1999). This wave of studies recognized that lifecycle variation should be accounted for, but assumed that the inclusion of age controls in the intergenerational regression would solve the issue. More recently, the literature has focused on non-classical measurement error and lifecycle bias. First discussed in Jenkins (1987), the problem gained attention with the generalized errors-in-variables (GEiV) model proposed by Haider and Solon (2006), and applications in different institutional contexts (e.g. Grawe (2006) for the US, Canada and Germany, Böhlmark and Lindquist (2006) for Sweden, Nilsen et al. (2012) for Norway). More recently, Nybom and Stuhler (2016) have shown that the GEiV model does not fully eliminate lifecycle bias. In parallel, Creedy (1988), Vogel (2007) and Hertz (2007) have suggested an alternative approach to deal with this problem, by modelling how the income profile varies with observable characteristics such as education or occupation. Related, Chau (2012) and Jäntti and Lindahl (2012) consider lifecycle models with heterogeneous intercepts and slopes.[3]

We contribute to this literature in two dimensions. First, we link it to the large literature on income processes (see Meghir and Pistaferri (2011) for a recent review). This link is interesting in both ways. One the one hand, a formal consideration of the income process and its key components allows for a more comprehensive understanding of the advantages and disadvantages of the previously proposed correction methods in the intergenerational literature, and informs the steps to be taken to better address the measurement problem. On the other hand, the intergenerational perspective contributes to the ongoing debate on the role of unobserved heterogeneity in the literature on income processes.[4] A controversial

[3]In light of the measurement problems in estimating the IGE, some recent studies focus instead on the intergenerational correlation in income *ranks*. A prominent example include Chetty et al. (2014). Correction methods for measurement error in ranks are proposed in Nybom and Stuhler (2017) and Kitagawa et al. (2019).

[4]We use the term "unobserved heterogeneity" to describe systematic variation in income growth that is not explained by standard individual characteristics, not in the more specific econometric sense.

question is if (residual) income grows at an individual-specific and deterministic rate (HIP) or follows a random walk (RIP). These two models are difficult do distinguish, and standard tests based on the covariance structure of income growth may not be very informative (Guvenen 2009). We argue that the type of intergenerational data that we use here is informative about this question, and provides direct and transparent evidence in favor of the HIP hypothesis – within educational or occupational groups, sons from high-income families have substantially steeper income growth than those from low-income families, in particular in the early stages of their career.

Second, and most importantly, we develop two new approaches to improve the IGE estimation in absence of complete income data for two generations. We aim to account for systematic differences in income growth between high- and low-income families, without trying to estimate fully heterogeneous slopes for each individual, which would yield noisy results in short income data. Although more research is needed for a thoroughly understanding of the data requirements for the adequate use of the proposed methods, our preliminary analysis suggest that our proposed lifecycle estimator performs fairly well in the most common scenario faced by practitioners: the one with a short panel with young sons and old fathers.

The proposed methods could be particularly useful for comparative work on mobility varaition across countries or over time. The observation that income inequality correlates positively with the intergenerational elasticity of income across countries (Blanden 2011, Corak 2013) and across regions within countries (Chetty et al. 2014) has received much attention. However, it is difficult to estimate the IGE in comparable ways in these settings. If the shape of income profiles differs across countries, a simple rule-of-thumb prescription to measure incomes around midlife is likely to introduce different biases in each country. This measurement problem has been particularly problematic for the estimation of mobility *trends*. Existing work is somewhat conflicting. Earlier work has found no evidence of shifts in mobility in the U.S. for the cohorts born between 1952-75 (e.g. Hertz 2007, Lee and Solon 2009), or later cohorts (Chetty et al. 2014), a finding that has received much attention given the large increase in income inequality over this period. However, more recent work argues differently. Davis and Mazumder (2019) show that the IGE declined sharply for cohorts born between 1957-64, in comparison with the ones born between 1942-53. Consistent with this, we demonstrate that prior methods do not reliably capture mobility trends in the Swedish registers, suggesting that lifecycle bias will be one important reason for these conflicting results. The explicit consideration of the income process may help to produce more comparable estimates.

The remaining sections of this paper are divided as following. In Section 3.2, we describe the Swedish registry data and our sample of analysis. In Section 3.3, we discuss the properties of income process and show evidence of each of its key components in the Swedish data. In Section 3.4, we analyze the current correction methods for the IGE estimates in light of the income process properties. In Section 3.5, we present our new correction methods and, in Section 3.6, we conclude.

3.2 Data

We use Swedish administrative panel data from the Institute for Evaluation of Labour Market and Education Policy (IFAU). The main population is comprised of individuals aged 16-64 at any point between years 1960-2013 (born 1896-1997). The income data consists of information on labor earnings. We use income data for the complete population for the years 1968, 1970, 1971, 1973, 1975, 1976, 1979, 1980, 1982 and 1985-2013. We also have a rich set of individual characteristics, such as education, occupation and family links. To abstract from female participation decisions, and as is standard practice in the intergenerational literature, we focus on father-child pairs in our estimation. Using data of the complete population, we first remove time effects. Specifically, we run a regression of log earnings of all individuals on year dummies that vary by educational group, controlling for age fixed effects. We then predict an earnings measure that is cleaned of the time fixed effects.[5] This is useful to obtain lifecycle profiles and mobility estimates that abstract from the business cycle conditions.

To observe nearly complete income trajectories for the offspring generation, we then restrict our sample to cohorts born in 1952-60. We only have continuous income data for the full population from 1968 onwards. We keep information for fathers born from 1927 on, such that we observe parental income between ages 41-65 for all fathers in our sample. Because young fathers are over-represented in our sample, our estimates are not representative for the Swedish population. This is not a concern for our analysis, as long as we observe a sufficiently heterogeneous sample to study lifecycle patterns in income. We impute data from the gap years with neighbor observations and we bottom code income in the first percentile of each cohort. Finally, we use all the data available on earnings to create a measure of lifetime income for both fathers and sons. For sons, we construct the lifetime income with data from 25-53 years old and for fathers with data from 41-57.

Our final sample of analysis is comprised of 201,063 father-son pairs, distributed across cohorts as shown if Table 3.1. Unsurprisingly, individuals belonging to younger cohorts are, on average, more educated, have more educated fathers and occupy higher-level position in

[5] We estimate the regression $Y_{igat} = \beta_g D_g + \beta_a D_a + \beta_{gt}(D_g \times D_t) + \varepsilon_{igat}$, where Y_{igat} is the log income of individual i, educational group g, age a and time t; D_g are dummy variables for the seven educational groups, D_a are age dummies, and $D_g \times D_t$ are dummy variables interacting all observable years (except the first year) with education. We then predict $Y_{igat} - \hat{\beta}_{gt}(D_g \times D_t)$, where β_{gt} is the group-specific time-slope.

Table 3.1: Descriptive Statistics of Main Sample

Cohorts	All	1952-54	1955-57	1958-60
Father-son pairs	201,063	38,265	67,738	95,060
Father's Age at Birth of Son	25.7	23.7	25.3	26.8
Log Lifetime Income of Sons	12.26	12.22	12.26	12.28
Percent of Zero Income Obs of Sons	7.2	6.7	7.0	7.5
Log Lifetime Income of Fathers	12.14	12.09	12.13	12.18
Percent of Zero Income Obs of Fathers	5.0	5.7	5.5	4.3
Education Level of Sons				
Primary	21.9	27.5	22.8	19.0
Short High School	39.3	35.6	39.2	40.9
High School	24.7	24.8	24.0	25.2
College	14.1	12.1	14.0	14.9
Occupation level of Sons				
Less than post-sec required	46.4	47.4	46.8	45.8
Post-secondary required	16.5	17.0	16.3	16.4
Professionals	20.2	19.2	20.2	20.6
Managers	17.0	16.3	16.7	17.3
Number of Missings	67,111	13,950	22,776	30,384
Education Level of Fathers				
Primary	52.7	58.3	53.6	49.7
Short High School	20.9	19.5	20.8	21.5
High School	18.8	17.0	18.3	19.9
College	7.7	5.2	7.3	9.0
Number of Missings	11,677	1,923	3,863	5,891

Notes: Log lifetime income variables abstract from time-fixed effects, as described in Section 3.2.

the labor market. This translates into a slightly higher log lifetime income of both sons and fathers.

3.3 The Income Process: Swedish and Simulated Data

Income fluctuations over the life cycle are the primary source of bias in intergenerational estimates. We illustrate here some properties of the income process that appear important in the intergenerational context. Our evidence will be based on actual income series as observed in the Swedish data, and on income series from simulated data calibrated to the U.S. economy. This evidence will then help us to characterize the advantages and limitations of the different correction methods (next section), and to develop those methods further (Section 3.5).

A large literature on income processes has studied the shape of income profiles over the life cycle, and decomposed its most important properties parametrically.[6] While the main properties of income processes are well established, two contrasting viewpoints exist about the idiosyncratic components of income growth over the life cycle. The *restricted income profile* (RIP) model views income as the sum of a mean-reverting component, which reflects "transitory shocks" that vanish over age, and a random-walk component, which reflects "permanent shocks" (see MaCurdy 1982). In contrast, the *heterogenous income profile* (HIP) model assumes that individual income is not subject to permanent shocks, but instead grows at an individual-specific and deterministic rate (see Guvenen 2009).

Following Guvenen (2009), let log income for individual i with experience h at time t be given by

$$Y^i_{h,t} = g(\theta_t, X^i_{h,t}) + y^i_{h,t} \tag{3.1}$$

where $g(\theta_t, X^i_{h,t})$ captures the part of variation that is explained by observable characteristics,[7] and

$$y^i_{h,t} = \alpha^i + \beta^i h + z^i_{h,t} + \phi_t \varepsilon^i_{h,t} \tag{3.2}$$

is the individual-specific income unexplained by those characteristics, where $z^i_{h,t} = \rho z^i_{h,t-1} + \pi_t \eta^i_{h,t}$ is a persistent and $\varepsilon^i_{h,t}$ a transitory shock. Persistent and transitory shock components are scaled by time-specific coefficients, as there have been important changes of those components over time (see Moffitt and Gottschalk 1995).[8]

The income process described by Guvenen (2009) has three key important components for our setting: (i) $g(\theta_t, X^i_{h,t})$, the income growth explained by observed characteristics; (ii) $\varepsilon^i_{h,t}$, the transitory noise; and (iii) $\alpha^i + \beta^i h$, the unexplained income growth.[9]

In Figure 3.1, we show evidence for the importance of each of these components in the Swedish data. Panel A illustrates that individuals' income profiles vary systematically with

[6]This evidence has been used explicitly in studies of the causal effect of parental income (e.g. Carneiro et al. (2015), but used mostly for motivational purposes in descriptive studies (an exception is Heidrich 2016).

[7]In Guvenen (2009)'s baseline model, the observable characteristic considered is a cubic polynomial in experience h. Yet, more generally, we could think as X^i_t as observables that could include education, gender, age, etc. The vector θ_t are the estimated coefficients, which are common to all individuals.

[8]Additionally, let $\varepsilon \sim (0, \sigma^2_\varepsilon)$, $\eta \sim (0, \sigma^2_\eta)$, $z_{0,t} = 0$ and

$$\begin{pmatrix} \alpha^i \\ \beta^i \end{pmatrix} \sim [0, \begin{pmatrix} \sigma^2_\alpha & \sigma_{\alpha\beta} \\ \sigma_{\alpha\beta} & \sigma^2_\beta \end{pmatrix}]$$

[9]Note that observed characteristics are the ones available in the data used by the practitioner, which include, in the most common scenarios, education or occupation. The unexplained income growth, instead, reflects variables that are unobserved by the researcher. These variables could include, for example, parental education, occupation or lifetime income. In the Swedish data, however, we do observe such characteristics and, therefore, can test whether these usually unobserved variables are important predictors of lifetime income.

Figure 3.1: Components of the Income Process

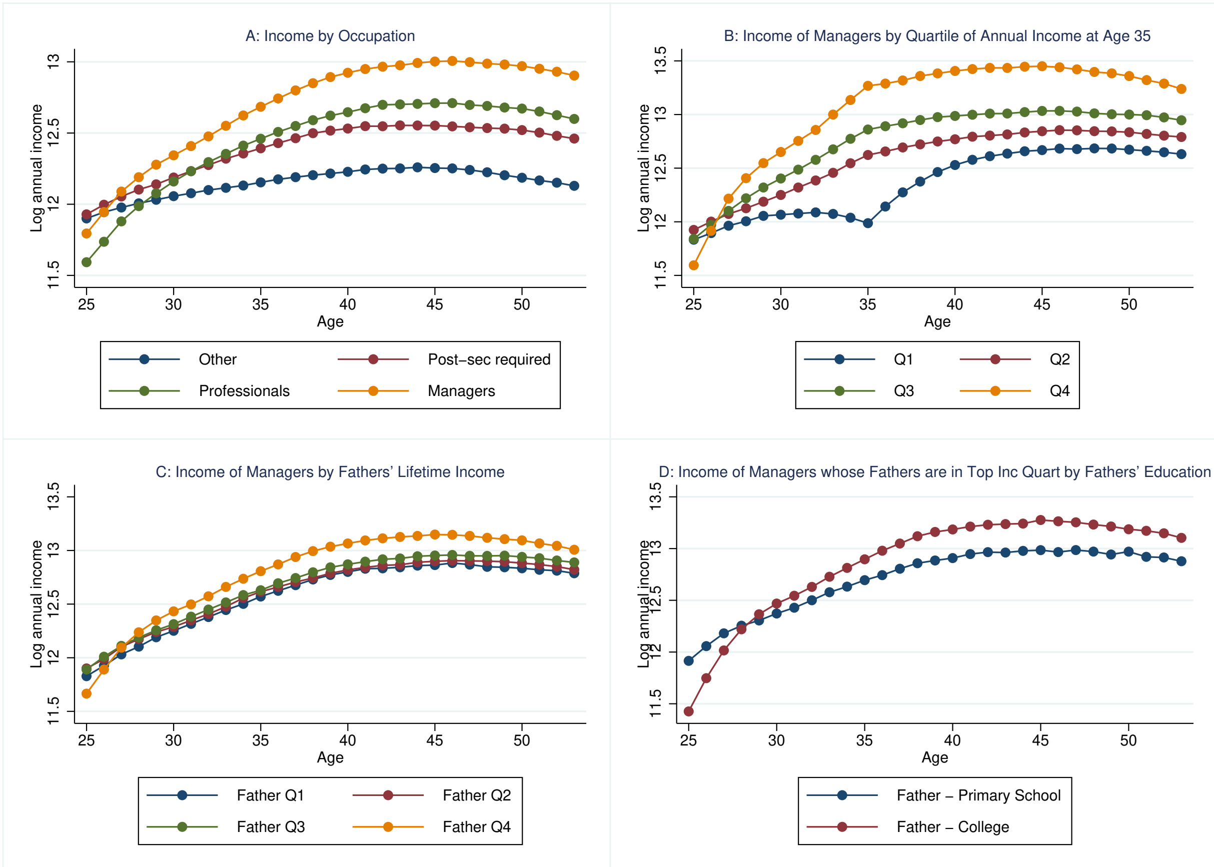

Notes: Panel A shows income trajectories by type of occupation. Panel B focus only on managers, who are split in four groups, according to their annual income at age 35. Category Q1 refers to the bottom quartile and Q4 to the top. In Panel C, managers are divided in four groups, according to fathers' lifetime income. Finally, in Panel D, managers whose fathers belong to the top quartile of lifetime income are divided in two additional groups: college-educated fathers and fathers with only primary school.

observable characteristics, such as occupation. Accounting for observable heterogeneity in gender, age, education or occupation is, therefore, important. Panel B illustrates that the level of income and subsequent income growth are negatively correlated. Specifically, we split individuals that work in a manager position into four quartiles of their annual income at age 35. Individuals in the bottom quartile have substantially stronger income growth in the following years, and individuals in the top quartile experience the lowest growth. Transitory shocks can rationalize this pattern, which matters because researchers often try to impute individual lifecycle profiles based on income levels at a particular age.

Panel C provides evidence on a more controversial question, namely if residual income $y^i_{h,t}$ grows at an individual-specific and deterministic rate (HIP) or follows a random walk (RIP). The panel plots the average income profiles of managers by the quartile of their *fathers'* lifetime income. We find that within each occupational group, sons' with fathers in the top income quartile have substantially higher income growth, but lower initial incomes. Managers from high-income fathers tend to have lower income in their late 20s, but nearly 30 percent higher incomes around age 40 than managers from less well-off families (Figure 3.1, Panel C). We find similar evidence within educational groups (Figure A.1, Panel C), or when considering other dimensions of family background. For example, in Panel D we consider managers whose fathers' income is within the top quartile in their generation, and plot their mean income trajectory by fathers' education. We find that among managers with high-income fathers, those with more educated fathers have steeper income profiles than those with less educated fathers.

The role of unobserved heterogeneity has remained controversial in the literature on income processes, because empirically it is difficult to distinguish from stochastic processes with high persistence. However, a consideration of family background yields direct, visual evidence on this question. Income growth varies systematically with parental characteristics, even after controlling for an individuals' own characteristics, such as gender, education, occupation, age and cohort. Because parental characteristics are predetermined with respect to, and potentially observed by the child, this pattern is more readily interpreted as a deterministic factor instead of a stochastic shock that would come as a surprise to the individual. As Guvenen (2009), we find that initial incomes and income growth are negatively correlated. Those with the highest lifetime income have the lowest initial incomes, and much steeper growth at early age. However, the individual-specific component of income growth is assumed to be linear in age in Guvenen (2009), while Panel A in (Figure 3.1) suggests that, in the Swedish context, unobserved heterogeneity matters most at young age, in the 20s and early 30s.

The fact that children from high-status parents tend to have steeper income profiles needs to be taken into account when estimating the intergenerational mobility of income. Researchers commonly predict income growth based on the observable characteristics of a child, such as his or her education and occupation. Panels C and D of Figure 3.1 suggest that accounting only for child's characteristics will be insufficient. To avoid a systematic bias with respect to parental income, parental characteristics need to be taken into account when predicting the likely paths of income growth among their children. We show below that failure to do so can lead to substantial bias in estimates of income mobility.

As an alternative way to evaluate different mobility estimators, we study their performance in simulated data. Motivated by the observation that income growth varies systematically with parental income in the Swedish registers, we simulate income profiles based on the

HIP process described in Guvenen (2009).[10] The process is calibrated to match moments of the income distribution in the United States, according to parameters estimated in Guvenen (2009) and reproduced in Appendix Table A.2. We use this calibrated process to simulate age-income profiles for a large set of individuals. Important for our purposes is how the various parameters of the income process correlate between parents and children, which does not follow from the existing literature. We therefore consider two different ways to model parental lifetime income for each simulated individual. In the first scenario, parental income is assumed to correlate with *all* sources of lifetime income, such that log parental lifetime income is simply a linear projection of the child's own log lifetime income. In the second scenario, parental income correlates only with the deterministic and individual-specific slope of income growth over the life cycle. We focus on evidence from the first scenario here.

3.4 Correction Methods in the Intergenerational Literature

The correction methods proposed in the literature can be classified into two alternative approaches, which vary in the specificity in which the income process is being considered: (i) errors-in-variables models that formalize the relation between (observed) annual and (typically unobserved) lifetime income, and (ii) models of the income process that indirectly determine this relation. Table A.1 lists recent applications and the correction methods on which they are based on. Most applications are motivated by the generalized errors-in-variables model proposed by Haider and Solon (2006), which yields the simple rule-of-thumb to measure incomes around midlife. Other studies instead predict individual-specific income profiles based on observable characteristics, as proposed by Creedy (1988), Vogel (2007) and Hertz (2007). We review these approaches in this section, and argue that they reduce but do not eliminate lifecycle bias from intergenerational mobility estimates.

3.4.1 Modelling Errors-in-Variables

Most work in the literature addresses the measurement problem by means of an errors-in-variables model that links (observed) short-run to (unobserved) lifetime income. Properties of the income process are not explicitly considered, but inform the errors-in-variables assumptions. These assumptions can therefore be interpreted as a reduced-form representation

[10] While our evidence supports the HIP model, other evidence speaks in favor of the RIP model (Hryshko (2012)). The distinction is important for the income process literature, but many of our arguments relate to properties of the income process that are common to both to models.

of those properties of the income process that are most important for mobility studies.

The Generalized Errors-in-Variables Model

In the classical errors-in-variables model as considered by Atkinson (1980) or Solon (1999), inconsistencies in the IGE are limited to attenuation bias caused by measurement error in the lifetime income of fathers. Jenkins (1987), Haider and Solon (2006) and Grawe (2006), however, show that the association between current and lifetime income varies systematically over the life cycle, contrary to a classical errors-in-variables model in which the errors are independent of true values. Haider and Solon (2006) therefore extend the classical into a *generalized errors-in-variables* (GEiV) model, in which this association is formalized.[11] Focusing on left-hand side measurement error, we estimate the linear projection

$$y_{sit} = \lambda_{st} y^*_{si} + u_{sit}, \tag{3.3}$$

in which y_{sit} is the son's of family i short-run or annual (log) income at age t, y^*_{si} is son's (log) lifetime income, and y^*_{si} and u_{sit} are uncorrelated by construction. Estimated at each age t, we obtain a series of $\hat{\lambda}_{st}$. We then estimate

$$y_{sit} = \beta_t y^*_{fi} + \varepsilon_{sit}, \tag{3.4}$$

where y^*_{fi} is parental lifetime income, obtaining a series of $\hat{\beta}_t$. Under the assumption that $Corr(y^*_{fi}, u_{sit}) = 0$ we have

$$plim \hat{\beta}_t = \frac{Cov(y_{sit}, y^*_{fi})}{Var(y^*_{fi})} = \beta\lambda_{st}, \tag{3.5}$$

where β is the true IGE based on measures of long-run income for both parents and children.

A key implication from equation (3.5) is that the use of short income spans in the child generation would not bias IGE estimates, if measured at an age in which λ_{st} is close to one.[12] Haider and Solon (2006) note that λ_{st} tends to be around one in midlife, such that estimates will be less biased when the data is drawn from that age range. This simple generalization of the classical error-in-variables model captures the relation between annual and lifetime incomes remarkably well. As shown in Table 3.2, the insight that λ_{st} will increase over

[11] The GEiV model has been extended further in subsequent work. Lee and Solon (2009) adapt it for the study of mobility trends. An et al. (2017) implement it within a non-parametric framework that allows for the IGE to be heterogeneous.

[12] The underlying intuition is that for individuals with different income trajectories there will nevertheless exist an age t^* at which the expected difference between individuals' (log) annual incomes equals the expected difference between their lifetime incomes.

age and be close to one around mid-age holds both in registry data from Sweden (see also Böhlmark and Lindquist 2006) and in simulated income series calibrated to the U.S. labor market.

Table 3.2: Lifecycle Bias and the Generalized-Errors-in-Variables Model

Sweden			U.S. (simulated)		
Son's Age	λ_{st}	β_t	Son's Age	λ_{st}	β_t
31	0.810	0.208	41	0.896	0.461
32	0.869	0.224	42	0.958	0.470
33	0.940	0.243	43	0.997	0.506
34	1.007	0.258	44	1.036	0.518
35	1.072	0.273	45	1.047	0.525
True		0.254	True		0.497

Notes: Estimates of λ_{st} and β_t are based on equation (3.3) and equation (3.4), respectively. The income process for the U.S. is based on Guvenen (2009).

However, the pattern reported in Table 3.2 also illustrates two problems with this approach. First, lifecycle bias may not be fully eliminated at the age at which λ_{st} is equal to one, because the shape of income profiles varies systematically with parental background (see Figure 3.1) in ways that are not fully captured by the son's own lifetime income. The GEiV model's underlying assumption that $Corr(y^*_{fi}, u_{sit}) = 0$ can therefore be violated (Nybom and Stuhler (2016)), and even exact knowledge of the lifecycle pattern of λ_{st} would not eliminate lifecycle bias.[13]

Second, and more serious in practice, is that the age at which $\lambda_{s,t} \approx 1$ is typically not known, as its estimation requires data on lifetime incomes. The applications listed in Table A.1 instead use the errors-in-variables model to motivate the rule-of-thumb to measure income at some point in midlife. But Haider and Solon (2006) warn that the optimal age is likely to vary across countries, and Table 3.2 demonstrates that even slight deviations from the optimal age t^* yield substantially different estimates of the IGE. In the Swedish data, estimates of the IGE vary by 29 percent (from 0.21 to 0.27) within the five-year window between age 31 and 35. This issue is not specific to the Swedish source. In simulated data based on moments of U.S. labor earnings (see Section 3.3), the IGE varies by about 15 percent (from 0.46 to 0.53) within a five-year window centered at the optimal age. As reported in Table A.1, the age at measurement varies considerably across studies. Estimates from the current literature are therefore likely to contain large biases, with unknown direction.

[13] This problem can be partially addressed by extending the generalized errors-in-variables model with covariates, see Nybom and Stuhler (2016).

The Generalized Errors-in-Variables Model and Mobility Trends

Lee and Solon (2009) provide an extension of the generalized errors-in-variables model for the estimation of mobility trends. They find that the IGE has remained quite stable in the U.S. income data during the 1980s and 1990s, a result that has received considerable attention, given the rise in income inequality over the same period. An important question is whether the issues discussed above only affect the estimated level of the IGE, or also its estimated trend over time. While the GEiV model may not fully eliminate lifecycle bias in the estimated level of the IGE, this bias would not affect the estimation of mobility trends over birth cohorts (or years) if that bias remains stable (i.e. if λ_{st} and $Corr(y^*_{fi}, u_{sit})$ remain constant over cohorts).

Figure 3.2: Trends in Lambdas and IGE Estimates

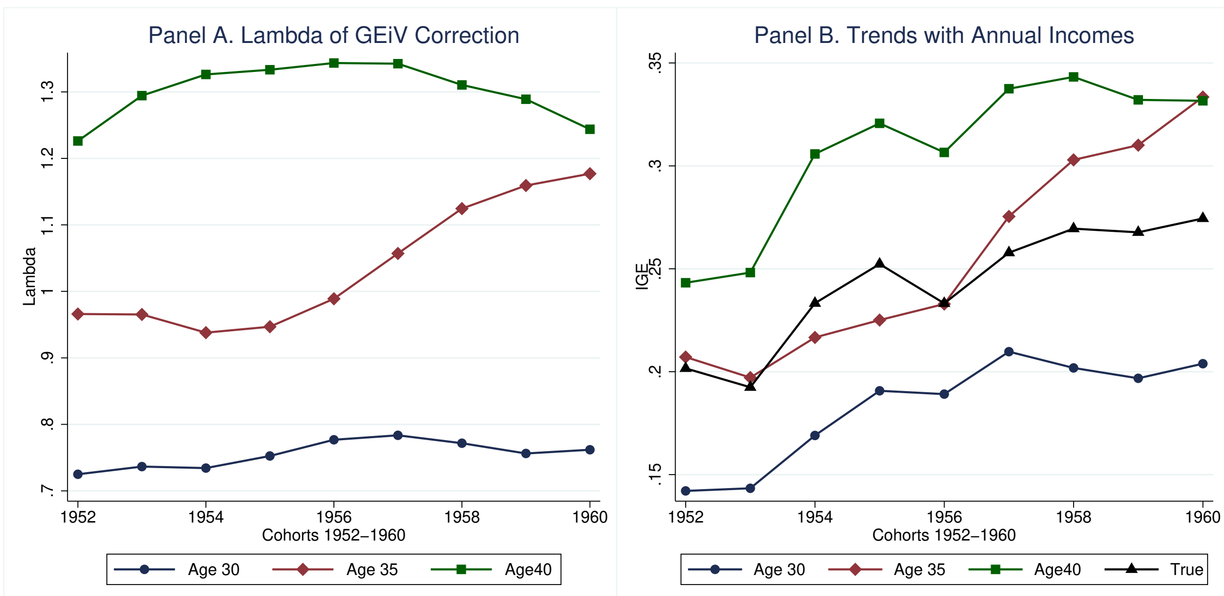

Notes: Estimates of λ_{st} (Panel A) and β_t (Panel B) are based on equation (3.3) and equation (3.4), respectively. Child income is based on a 3-year moving average centered around the cohort indicated on the x-axis.

However, the structure of income profiles has changed over time in the U.S. (Guvenen, 2009) and other countries, suggesting that the age-profile of λ_{st} may also change over cohorts. Panel A of Figure 3.2 shows that it indeed shifted over the cohorts contained in our Swedish sample. Measured at age 35, estimates of λ vary between 1 and 1.2 over a cohort range of less than one decade. According to equation (3.5) , estimates of the IGE based on short-run income will vary by a similar amount. These magnitudes are large compared to the type of mobility shifts that we may expect over time. To illustrate this point, Figure 3.2 Panel

B plots the "true" mobility trend in our Swedish sample (as based on long series of income for both parent and child generations) as well as the estimated trend when son's income is only observed at a particular age (at age 30, 35 or 40, respectively). The estimates based on annual data capture the direction of the trend, but misstate its magnitude. The IGE increases over our observation period (black line), but it appears to increase twice as much when child income is measured at age 35 (red line). This observation reflects the fact that income distributions, and therefore the value of λ, can change substantially with macroeconomic conditions (such as the recession that Sweden experienced in the early 1990s).[14] We therefore argue that mobility estimates based on short-run income measured around a fixed age will not be a good approximation of gradual trends in mobility over time (but may still identify sudden and large shifts in mobility).

3.4.2 Modelling the Income Process

An alternative is to directly model the income process. After estimating individual profiles based on observable characteristics and income, we can predict an income measure for each person, and use in the intergenerational estimation. The key issue is that typically only short income spans are available, so their extrapolation to the complete life cycle requires parametric assumptions on the shape of age-income profiles.

We review methods by Hertz (2007), Vogel (2007), and Creedy (1988), who propose different solutions to this problem. They capture important properties of the income process, but neither attempts to capture each of the three key properties highlighted in Figure 3.1. The approach by Creedy (1988) accounts for both observable and unobservable determinants of individual-specific income growth, but does not account for the influence of transitory income shocks. The approaches by Vogel (2007) and Hertz (2007) allow for unobserved individual heterogeneity in income *levels* (i.e. individual fixed effects) and observed individual heterogeneity in income *growth* as explained by own characteristics (such as education or occupation). However, even conditional on those characteristics, children from high-income parents tend to have steeper income growth (see Panel C in Figure 3.1).

The resulting estimates of mobility levels or trends therefore still contain some lifecycle bias. Specifically, our evidence suggests that previous work may have understated the degree to which the intergenerational elasticity has increased in the United States. This implication is significant, as an increase in the elasticity should have been expected based on theoretical

[14] In the absence of direct data on lifetime incomes, it is less clear if this observation extends to the US case. Lee and Solon (2009) note the rise in income inequality over their sample period, but argue that the signal-to-noise ratio in log current income as a proxy for log long-run income may have remained stable (according to the estimates in Gottschalk and Moffitt (1994), and Haider (2001)). In contrast, Hertz (2007) argues that age-income profiles have not remained stable.

arguments, and based on the substantial increase in the variance of income in the child generation. Our results imply here that methodological improvements might be required to make a more definitive statement on the time trend of the intergenerational elasticity in the U.S. or other countries.

Observable Heterogeneity with Individual Fixed Effects

A common approach is to model the income process as a function of age, observable characteristics, and individual fixed effects. Hertz (2007) and Vogel (2007) propose slightly different implementations of this approach, but their starting point is similar. In a first step, we estimate income profiles with individual fixed effects and growth rates that depend on education (and possibly other variables). For example, the age-correction in Hertz (2007) is based on the equation:

$$log(Y_{ict}) = \alpha_i + A_{ict}\beta + A_{ict}Z_{ic}\gamma + \varepsilon_{ict}, \tag{3.6}$$

where t is the period of observation, c is the year of birth, α_i are individual fixed effects; A_{ict} represents a quadratic equation on age; and $A_{ict}Z_{ic}$ represents the interaction of each of the age variables with a vector of predictors of the shape of the age-earnings profile.[15] The corresponding first-step equation in Vogel (2007) is similar.[16] Because Hertz' primary objective is to estimate mobility trends, this equation is estimated separately for each cohort in the sample, relaxing the assumption of a constant age-income profile across cohorts.

However, intergenerational studies tend to observe children at young age, while parents are observed only at older age. It is therefore not straightforward how to estimate complete lifecycle profiles based on equations (3.6) or (3.7). Hertz (2007) and Vogel (2007) address this problem in different ways. Hertz predicts the income for each person *at one particular age*, such that incomplete profiles do not need to be extrapolated over the entire age range. Vogel instead aggregates incomes *over the complete life cycle*. To address that only part of the life cycle is observed in either generation, he assumes that the shape of age-income profiles remains similar in the two generations. One can then pool the income information for parents and children for estimation of the first-step equation (3.7), predict income $log(\hat{Y}_{ict})$ at each age, and construct measures of lifetime income for each individual. The predicted

[15] Hertz stresses a number of advantages of this two-step procedure as compared to the errors-in-variables approach as implemented in Haider and Solon (2006) and Lee and Solon (2009); (i) it permits the inclusion of individual fixed effects for the age-income profile, and (ii) it takes into account other observable sources of heterogeneity, such as education. If education affects the shape of age-income profile even after conditioning on parents' income, this omission is a source of inconsistency.

[16] Vogel (2007) estimates

$$log(Y_{ict}) = \alpha_i + A_{ict}\beta + \gamma t + \varepsilon_{ict}, \tag{3.7}$$

separately for each education group. The main differences are the inclusion of a quartic instead of quadratic in age, and the inclusion of a linear time trend (vs. separate estimations by cohort.)

lifetime incomes (Vogel) or age-specific incomes (Hertz) are then used for estimation of the IGE in a second step.

By modelling important determinants of the income process such as education or occupation (see Panel A in Figures 3.1 and A.1), these methods will reduce lifecycle bias in mobility estimates. However, income growth varies also within education or occupation groups, and these differences are also systematically related to parental lifetime income (see Panel C in Figures 3.1 and A.1). We show next to what degree this unobserved heterogeneity affects estimates of *level* and the *cohort trend* of the intergenerational elasticity. We illustrate our arguments based on actual income series from the Swedish sample, but they likewise hold in simulated income series drawn from the income process described in Guvenen (2009).

Estimating Levels

In Table 3.3 Panel A, we report how the approach by Hertz (2007) performs in the Swedish sample. The row "True" reports our baseline estimate for the IGE. The row "Annual" reports the IGE estimate when sons' income are measured at age 25, 30 or 35. Measurement at a too early age leads to large downward biases, while estimation at a late age generates an upward lifecycle bias. The other rows then report the IGE estimate based on different implementations of equation (3.6), varying (i) the sampling range for this first-step estimation (from 25-30 to the full 25-53 range) and (ii) the age at which incomes are predicted for the second-step estimation (age 25, 30 or 35).

Application of Hertz' method reduces the lifecycle bias compared to raw annual estimates.[17] However, estimates based on this method still exhibit a clear lifecycle pattern, increasing systematically with respect to (i) the age range on which the first-step equation (3.6) is estimated, and (ii) the age at which individual incomes are then predicted. For example, the estimate shrinks from 0.167 to 0.143 or 0.071 when the sampling range shrinks from the full range to age 25-40 or 25-30, respectively.

The intuition here is that the estimates of the fixed effects depend on the age range included in the profile estimation. In particular, when observing only early ages, we are understating the lifetime income of those with low initial incomes but steep profiles. The issue is illustrated in Figure 3.3 Panel A. Suppose the blue line is the true income trajectory of individual i with a steeper than average profile, while the red line is the mean estimated income, before adjusting with the individual fixed effect. Now, suppose we only observe sons from 25-40. In this case, the green line will represent the adjusted estimated income for individual i, which is nothing more than the red line plus a negative individual fixed effect.

[17] The reason is that the fixed effects in the first-step prediction can partially capture the substantial variation of incomes *within* education groups.

Table 3.3: Lifecycle Estimators in the Swedish Data

Observed Range	Panel A: Predict at Age			Panel B: Predict
	25	*30*	*35*	Complete Profile
25-53	0.197	0.235	0.254	-
25-45	0.167	0.206	0.234	-
25-40	0.143	0.192	0.234	0.255
25-35	0.109	0.175	0.241	0.235
25-30	0.071	0.181	-	0.218
True	0.254	0.254	0.254	0.254
Annual	0.030	0.193	0.273	-

Notes: This table reports estimates from the methods proposed by Hertz (2007) and Vogel (2007) in the Swedish data. In Panel A, we implement the method proposed by Hertz (2007) . The observed range refers to the age range of annual incomes used in the first step (equation 3.6 or (3.7). Age of prediction refers to the age at which income is predicted in the second step (e.g. age 25 in Hertz 2007). In Panel B, we implement the method proposed by Vogel (2007), in which incomes are instead predicted over the entire lifecycle.

Due to the negative fixed effect, we understate the the lifetime income of those with steeper profiles, and therefore, the intergenerational elasticity. The shorter and earlier the age range, the more we are understating the elasticity, as we can see from Table 3.3.

A similar issue affects the approach proposed by Vogel (2007). Again the approach understates the lifetime income of sons (observed early in life) with steeper than average profiles, and overstate the lifetime income of fathers with steep profiles (observed late in life). This argument is illustrated in Figure 3.3 Panel B. In red, we have the estimated average profile from one educational group and, in blue, we have an example of the true income for individuals with steeper profiles. Suppose we observe sons only earlier in life (ages 25-40 in the picture) and fathers only later in life (ages 40-55). Such individual heterogeneity can only be captured by the fixed effects, so the father will have a positive fixed effect and the son a negative fixed effect. As a consequence, we would be understating the lifetime income of sons (green line), overstating the lifetime income of fathers (gray line), and therefore, understating the intergenerational elasticity. This argument is illustrated in Table 3.3. As explained conceptually, the estimate of the IGE understates the true value when using income observations at a young age (e.g. 25-30 and 25-35), and overstates the true value when using observations at an old age (e.g. after age 40).

Estimating Trends

As explained in Hertz (2007), such lifecycle bias may be less problematic if our objective is to understand the *trend* in mobility instead of its level. As long as the remaining lifecycle

Figure 3.3: Illustration of Potential Problems with Fixed Effect Estimators

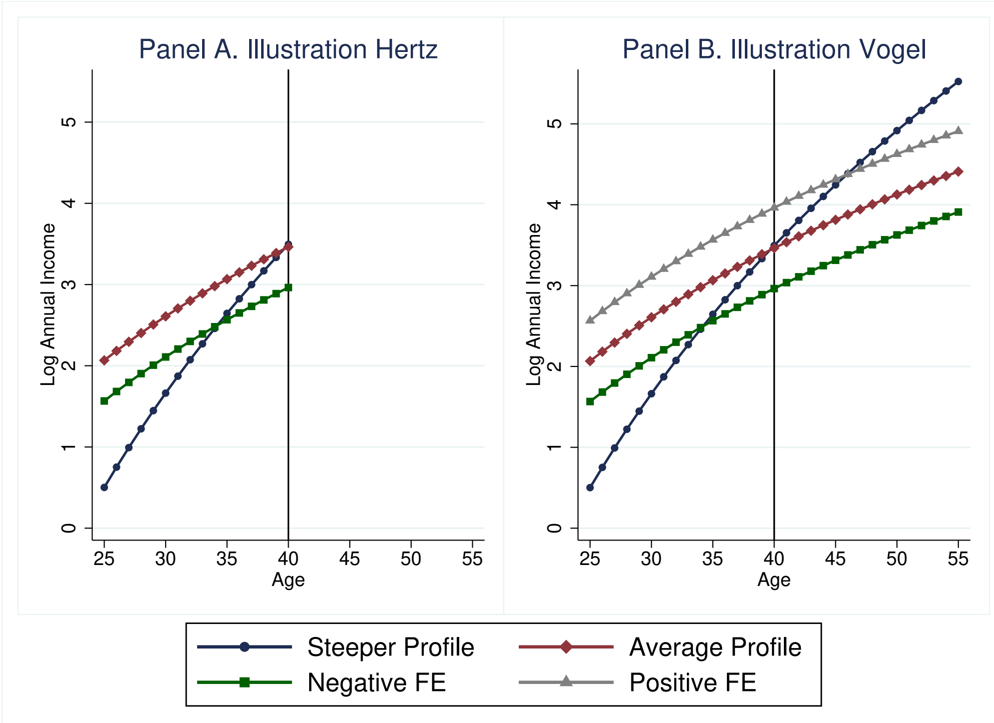

Notes: In the Figure, the blue line represents the income trajectory of individuals with steeper than average profiles, while the red line represents the estimated average income profile. The green line is the average income adjusted by a negative individual fixed effects, while the gray line is the adjustment by a positive fixed effect.

bias remains similar across cohorts, we may correctly capture that trend. Hertz keeps the age at which incomes are predicted constant over cohorts, eliminating variation of the IGE estimates across the horizontal dimension of Table 3.3. However, estimates of the IGE vary also with the age range for estimation of the first-step equation, because the estimates of the fixed effects depend on the age range included in the profile estimation (as illustrated in Figure 3.3). The shorter and earlier the age range, the more we are understating the IGE, as we can see from Table 3.3.

This observation is important, because applications rarely hold the age range of the sample constant when estimating mobility trends (given the limited availability of data for the more recent cohorts). For example, the younger cohorts in Hertz (2007) are estimated on a shorter and earlier age range (see Table 2 in Hertz 2007), and suffer then from a stronger downward bias. To illustrate the potential magnitude of this problem, we replicate it in the Swedish sample in Figure 3.4. In black, we plot the "true" trend based on our baseline estimate, which suggests that the IGE has been increasing in our sampling period. In red, we plot estimates based on predicted incomes at age 25, as based on the first-step estimation of equation (3.6) in income from ages 25-45 (blue line) or age 25-35 (red line). While their levels are off, these estimates do replicate the increasing trend in IGE estimates. However, these results were based on a fixed sampling range, while the sampling range varies in many applications. If we reduce the age range for the more recent cohorts, the increase in the trend is substantially understated (green line). Our findings therefore suggest that the preferred estimates in Hertz (2007) and similar applications understate the cohort-trend in the IGE in the United States.

Figure 3.4: Estimation of Trends in the IGE

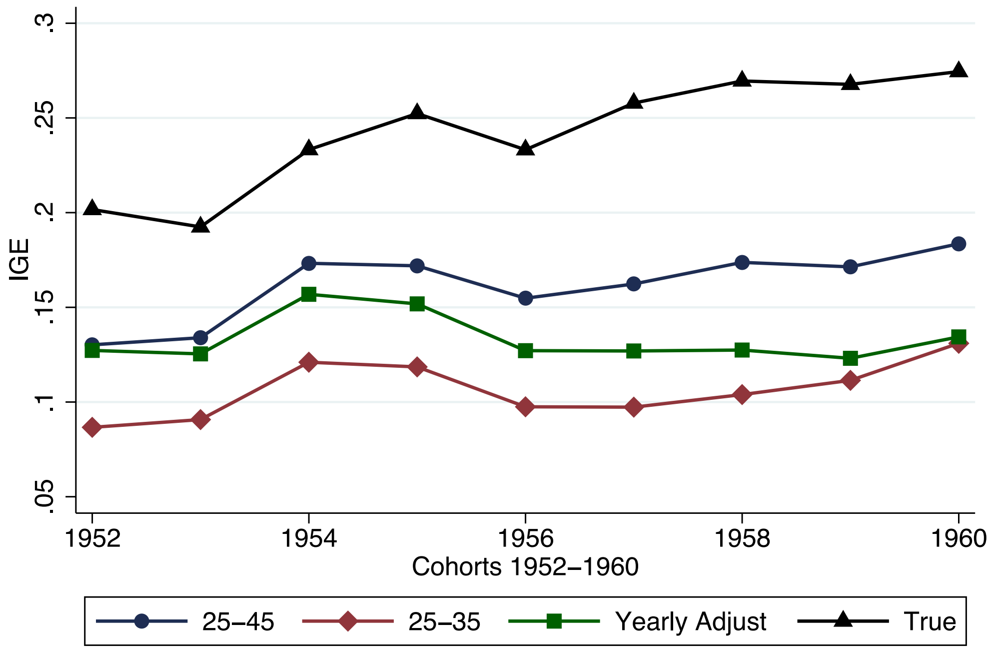

Notes: In this Figure, we plot trends in the IGE using the Swedish data. In black, we plot the true IGE. In blue and in red, we plot the IGE corrected by the Hertz method using, respectively, ages 25-45 and 25-35 as the age-ranges for estimation. Finally, in green, we plot the yearly adjusted estimates, similar to the one implemented in Hertz (2007). Here, the age range for estimation reduces as the cohorts become younger. For the 1952 cohort, the age range for estimation is 25-43, for the 1953 cohort, the range is 25-42, and so on. For the 1960, the youngest cohort, the age range is 25-35. As in Hertz (2007), the age used for prediction is 25.

Extrapolating from Observable Profiles

An alternative approach is to adjust annual incomes based on information on average age-earnings profiles obtained from external data sources (e.g. Atkinson 1980, Jenkins 1987). Creedy (1988) expands on this idea, developing a correction method based on the insight that the dispersion of earnings tends to increase over age, even conditional on education or occupation (see Figure A.2). Instead of estimating a separate income profile for each individual, he assumes that growth rate depends on the *rank* of the individual in the income distribution. An important advantage of this method is that it can be implemented in cross-sectional data sources.

In a first step, we estimate how the mean and the variance of log income vary over age within each occupational or educational group. Following Creedy (1988), we estimate

$$log(Y_{ij}) = \beta_0 + \beta_1 age_{ij} + \beta_2 age_{ij}^2 + u_{ij}, \tag{3.8}$$

separately by each occupational or educational group j, where Y_{ij} is the income of individual i and group j. Then, we predict $\hat{\mu}_{tj}$, which is the average income by each occupational group j and age group t. The variance of log income σ_{tj}^2 is also computed within each group. Then, we estimate:

$$\sigma_{tj}^2 = \beta_0 + \beta_1 age_{tj} + \epsilon_{tj}, \tag{3.9}$$

and obtain predicted values for $\hat{\sigma}_{tj}^2$. Alternatively, one can obtain these measures from external sources.

In a second step, these predicted values are used to rescale individual incomes to a common base year. First, compute the *standardized* value of an individual's log-earnings,

$$z_t = log(Y_t) - \hat{\mu}_{tj}/\hat{\sigma}_{tj}. \tag{3.10}$$

Then, rescale these standardized incomes according to the occupation or education-specific age-earning profile to compute adjusted log earnings at a common age $t*$:

$$log(Y_{t^*}) = \hat{\mu}_{t^*j} + z\hat{\sigma}_{t^*j}. \tag{3.11}$$

Those adjusted earnings depend on a single observable income at age t and on the values of $\hat{\mu}_t$ and $\hat{\sigma}_t$ that were predicted within the educational and/or occupational group. Finally, we have adjusted income observations for different ages, computed based on a single cross-section observation and scaling factors. Creedy proposes to either use adjusted earnings directly or to compute an aggregated discounted lifetime earnings measure for the estimation of the IGE.

We implement this method in the Swedish data. We combine the first-step estimates of $(\hat{\mu}_{tj})$ and $\hat{\sigma}_{tj}^2$ with an individual's earning at a certain age, to obtain his predicted income from ages 25-53 (which are then used to construct lifetime incomes). We therefore obtain a different measure of lifetime income, and a different estimate for the IGE, depending on the age at which we measure sons' income. We plot the resulting estimates of the IGE in Panel A of Figure 3.5. We compare estimates based on the "true" lifetime income of sons (black line), estimates based on annual incomes (green line), and two versions of Creedy's proposed

Figure 3.5: Extrapolating from Observable Profiles

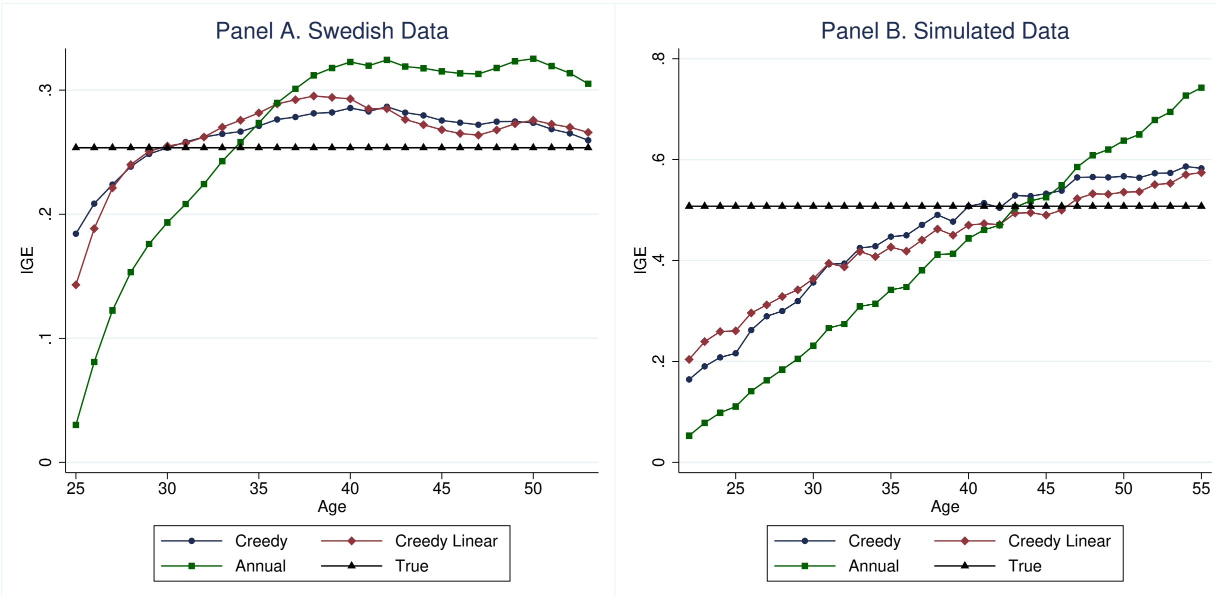

Notes: In the Figure, we compare the implementation of the Creedy method in the Swedish versus simulated data. The green line represents annual estimates, the blue line the estimates after Creedy's correction using true values of the mean and variance of income and the red line the Creedy implementation with linear estimates for the mean and the variance, according to equations 3.8 and 3.9.

estimator. In the first, we approximate the profiles of $(\hat{\mu}_{tj})$ and $\hat{\sigma}^2_{tj}$ with a linear function in age (red line). In the second, we use their non-parametric age profile as observed in the sample (blue line).

The comparison demonstrates that estimates of the IGE can be significantly improved by taking the dispersion of income growth over age into account. The corrected estimates are within 20 percent of the benchmark over the age range 30 to 50, even if the age profiles of $\hat{\mu}_{tj}$ and $\hat{\sigma}^2_{tj}$ are approximated linearly. The correction is not as good as in Haider and Solon (2006), but it is also based on substantially less ambitious data – only the age pattern of the variances and means is required. As Creedy discusses, the statistics that are necessary for the correction can potentially be estimated from a single cross-section. However, Figure 3.5 also shows that the correction method works only imperfectly, and tends to overstate the IGE over most of the age range. We study the performance of this method further in income series based on U.S. data (see Section 3.3), in which we can more exactly point to the components of the income process that can or cannot be accommodated by this method. Creedy's correction does not work well in this simulated data, as shown in Panel B of Figure 3.5.

A key limitation is that equation (3.11) rescales incomes based on the assumption that individual's rank in the widening income distribution remains stable over age. That is, it is assumed that individuals with high rank in annual incomes have higher income growth in the future. This is not the case in simulation or data, as illustrated in Panel B of 3.1. Because of the presence of short-term noise (i.e. the AR(1) process in the simulation), annual incomes are instead mean-reverting – individuals with high income rank at age t tend to have lower income growth in the next few years. In short, the method proposed by Creedy (1988) does not address for the presence of transitory noise. In the HIP process proposed by Guvenen (2009), for example, incomes at early ages are dominated by shocks and HIP intercepts, but incomes at late ages and log lifetime incomes are dominated by growth rates. As a result, the imputation in equation (3.11) overstates the variance of lifetime incomes, and thus tends to overstate the IGE (if applied to offspring on LHS of equation). The method performs performs better the more important the heterogeneous growth rates are compared to the transitory shock component.

3.5 New Correction Methods

In this section, we propose two distinct methods to address lifecycle bias in the estimation of the intergenerational elasticity. In the first method, we propose a *lifecycle estimator* that extends previous work by explicitly allowing for children from high-income parents to have steeper income profiles than their peers. In the second method, we consider a "standardized" errors-in-variables model, which relates the generalized errors-in-variables model proposed by Haider and Solon (2006) to moments that are more directly estimable in practice.

3.5.1 The Lifecycle Estimator

In this section, we consider estimation of the intergenerational elasticity by means of a "*lifecycle estimator*". In a first step, we estimate the shape of individual income profiles over the life cycle. This allows us to predict the *complete* lifecycle profile, and therefore lifetime income, for each person. In the second step, we estimate the intergenerational elasticity based on the predicted lifetime incomes. This *lifecycle estimator* uses the available income information more fully than the rule-of-thumb implementations based on income averages that are prevalent in the current literature. We illustrate that the estimator can be applied in a wider range of data scenarios, and that it provides more robust and more comparable estimates than the estimator based on income averages.

Our estimator builds on existing methods to predict lifecycle income profiles based on observable characteristics (Vogel 2007 and Hertz 2007) or income ranks (Creedy, 1988). As shown in Section 3.4, these estimators perform substantially better than naive estimators based on annual incomes, but still suffer from systematic lifecycle biases. No existing method addresses the three key components of the income process described in Section 3.3: (i) income growth explained by observable characteristics; (ii) transitory noise and (iii) unexplained income growth. By combining the relative strengths of previously proposed methods, we hope to account for all three components, and to address lifecycle bias more fully.

In a first step, we estimate individual lifecycle profiles. We allow for income profiles to vary with observable characteristics such as age and education, and allow for individual fixed effects (as Vogel 2007 and Hertz 2007). We however also account for the fact that children from high-income families tend to have steeper slopes, even conditional on their own observable characteristics (see Figure 3.1 in Section 3.3).[18] Specifically, we estimate

$$log(Y_{ict}) = \alpha_i + A_{ict}\beta + A_{ict}Z_{ic}\gamma + Age_{ict}P_{ic}\delta + \varepsilon_{ict}, \tag{3.12}$$

where t is the period of observation, c is the year of birth, α_i are individual fixed effects, A_{ict} represents a quadratic equation on age; $A_{ict}Z_{ic}$ represents a linear interaction of age with a vector of the individual's own characteristics (such as education); and $Age_{ict}P_{ic}$ represents a linear interaction of age with family characteristics or other predictors of the shape of the age-earnings profile among individuals with similar observable characteristics.

In a first specification (the *2-step parental* estimator), P_{ic} contains parental income and four indicators for parental education. This specification controls rather directly for variation in income growth by parental background. However, parental education and/or income will not always be observed in the data at hand. We therefore consider also alternative specifications that do not require the observation of parental characteristics.

In a second specification (the *2-step FE linear* estimator), we allow individual slopes to vary systematically with the level of an individual's own income. The underlying intuition is that individuals from high-income families have both higher levels and steeper slopes than those from low-income families (see Figure A.1, Panel C). In settings in which we cannot directly control for parental background, we may therefore still be able to capture some of the systematic heterogeneity in slopes by considering their co-variation with income levels (which are partially observed).[19] To implement this idea, we first estimate equation (3.12)

[18] Such heterogeneity can potentially be captured by (i) estimating separate slopes for each individual (see Jäntti and Lindahl 2012), or by (ii) estimating how individual slopes vary with parental background. We do not pursue the first option, as direct extrapolation from partially observed slopes may produce unstable predictions of lifetime income if only few income observations are available per person (as is typically the case in the intergenerational context). Instead, we allow for individual slopes to vary systematically with parental characteristics or an individual's own income history.

[19] This estimator is inspired by the observation of Creedy (1988) that income profiles tend to "fan out" over age, and that individuals with higher income levels tend to have higher income growth. By measuring

without the $Age_{ict}P_{ic}$ interaction, to yield estimates of the individual fixed effect $\hat{\alpha}_i$. We then re-estimate the equation with the $Age_{ict}P_{ic}$ interaction, setting $P_{ic} = \hat{\alpha}_i$.[20] Finally, we consider a variant of this strategy in which we interact the fixed effect with a quadratic instead of linear function in age (the *2-step FE quadratic* estimator). The choice of quadratic interaction is motivated by the observation that income growth varies more strongly in the early career than later in life.

Three conceptual issues arise in the estimation of lifecycle estimators, such as equation (3.12). First, the estimation consists of two or even three steps, which affects inference and the estimation of standard errors. However, sampling uncertainty is not a major concern in our setting, since our samples are very large. We therefore ignore uncertainty from the first stage. Second, the child generation is typically only observed at young age, so the shape of income profiles need to be extrapolated over the non-observed age range. This issue can be addressed in different ways, depending on data availability and context.[21] We therefore abstract from the issue here, exploiting the fact that we do observe long income profiles in our main sample. Specifically, we split each income profile in our child generation into two copies, with income for the "younger" copy assumed to be observed only up to some age threshold, and the "older" copy being observed thereafter. We thereby focus on the problem of missing income information for any given person, while abstracting from the issue that certain age ranges might not be observed for *any* comparable individual in the population.

A third conceptual issue is that the dependent variable in equation (3.12) is the *logarithm* of annual income. For the construction of (absolute) lifetime income, the predictions from this model need to be converted to absolute incomes, which gives rise to a well-known *re-transformation problem*: while the fitted values from the regression have mean zero by construction ($E[\hat{\varepsilon}_{ict}] = 0$), their mean will be positive after the transformation ($E[exp(\hat{\varepsilon}_{ict})] > 0$). To our knowledge, this issue has so far been ignored in the intergenerational literature. In itself, a non-zero mean is not problematic for intergenerational estimators. If the term is constant across individuals, it would be linearly separable in log lifetime income and therefore only affect the intercept of the intergenerational regression, not the intergenerational

an individual's position with the fixed effect α_i (instead of the individual's annual income or income rank) we however avoid the issue that annual incomes tend to be mean reverting because of transitory noise (see Section 3.3).

[20] In general, the individual fixed effects from this second regression will not coincide with the fixed effects as estimated from the first regression. To reduce these deviations we could estimate the equation multiple times, updating $P_{ic} = \hat{\alpha}_i$ after each regression. We do not implement this procedure here, as it appears to have only a negligible effect on estimates of the IGE.

[21] Potential solutions are to extrapolate based on functional form assumptions, to pool the child and parent generation based on the assumption that the shape of income profiles are sufficiently similar (as proposed by Vogel 2007), or to use an auxiliary data set in which a single generation is observed over a sufficiently long time period, or in which individuals of different ages are observed around the same time period.

elasticity. However, $exp(\hat{\varepsilon}_{ict})$ tends to have a larger mean for individuals with low lifetime income in our data, because their incomes tend to be more variable over the lifecycle. This systematic variation would bias estimates of the intergenerational elasticity upwards. Because the re-transformation problem is a generic problem that affects any estimator based on log annual incomes, we abstract from it when constructing our estimators.[22]

We present evidence on the performance of our proposed lifecycle estimators in Table 3.4.[23] To probe their performance in different data settings, we consider different age thresholds, assuming that child income is observed only between age 25 and 27 (first row), age 25 and 30 (second row), and so on. In column (1), we report our benchmark estimate (i.e. the direct estimate based on lifetime incomes for both child and parent generation), which is approximately $\hat{\beta} = 0.26$.[24] In column (2), we report an estimate based on annual incomes for the child generation. Specifically, we measure annual incomes at the median age within each age bracket (e.g. age 26 in the first row, and age 28 in the second row). As shown previously, the estimated elasticity is very sensitive to the age at measurement. It is as low as $\hat{\beta}_{annual} = 0.08$ when incomes are measured between age 25 and 27, and generally deviates substantially from the benchmark estimate.

In columns (3) to (6) we implement the lifecycle estimators. In column (3), we report a baseline lifecycle estimator that is based on estimating the first-step equation (3.12). We estimate this equation separately for four education groups (as defined in Figure A.1), instead of including education as a separate predictor. This estimator corresponds closely to the estimator proposed by Vogel (2007), and performs accordingly. While performing much better than the naive estimator based on annual incomes, it still deviates substantially from the benchmark estimate, and also varies systematically with age. The estimates are still as low as $\hat{\beta}_{baseline} = 0.17$ when incomes are measured between age 25 and 27.

In column (4), we report the *2-step parental* estimator, in which we allow for income growth to vary systematically with parental education and income. The estimator performs better than the baseline lifecycle estimator, in particular when incomes are only observed at very young ages. The estimates are $\hat{\beta}_{parental} = 0.21$ when incomes are measured between age 25 and 27, and are within 10 percent of the benchmark estimate when income at older ages are observed. In columns (5) and (6), we report the *2-step FE linear* and *2-step FE*

[22] Specifically, we estimate a complete life cycle profile of each individual in the child generation, based on a quartic in age and individual fixed effects, and then construct $SM_{ic} = \sum_{t=25}^{53} exp(\hat{\varepsilon}_{ict})$ as a measure of the individual income variability around the mean tendency. We use this adjustment factor to adjust predicted lifetime income for each of the lifecycle estimators. See Wooldridge (2006), for a discussion of the re-tranformation problem and its potential solutions.

[23] These estimates are based on a random 50% draw of our main sample. We will replace them with results from the full sample in the final manuscript version.

[24] Because we hold our sample constant across all estimators, and because because some individuals may not be observed in certain age ranges, the benchmark estimate varies slightly with the chosen age threshold.

quadratic estimators. They perform slightly better than the 2-step FE parental estimator. This finding suggests that systematic variation in income growth by parental background can be addressed without observing parental characteristics, simply by taking into account that individuals with higher income levels also tend to have higher income growth.

The 2-step FE quadratic estimator is therefore our preferred specification. Estimates from this approach are generally within 5 percent of the benchmark estimate, with a slight larger understatement of the intergenerational elasticity when incomes are observed only at very young ages far below age 30. The estimates are otherwise remarkably insensitive to the age at which incomes are measured, fluctuating around the benchmark without any apparent systematic lifecycle bias. This estimator therefore promises to yield precise estimates of the intergenerational elasticity, irrespectively of the type of income snapshots that are observed for the child generation. Moreover, it yields estimates that are quite comparable even when the underlying samples are observed for very different age ranges.[25]

Finally, we test how the performance of the lifecycle estimator varies with the number of income observations that are available for each individual in the child generation. In Table 3.5, we report estimates from the 2-step FE quadratic estimator in different data scenarios. Specifically, we randomly select 6 income observations for each person within the indicated age range, and then (randomly) drop income observations until only 2 annual incomes are observed per person. Reducing the number of income observations increases the noise in the estimation of equation (3.12), and the noise in the predicted lifetime incomes. The regression R2 therefore decreases when fewer income observations are available. However, estimates of the intergenerational elasticity remain remarkably stable. There is no clear pattern when incomes up to age 35 are observed. The estimates appear to decrease slightly when incomes are observed only up to age 30, presumably because the estimate generally tends to slightly understate the intergenerational elasticity when income is observed only at very young age (cf. Table 3.4).

Overall, the lifecycle estimator appears to perform very well. It can nearly eliminate lifecycle bias in our data, with estimates fluctuating closely around the benchmark. Moreover, the approach yields estimates that appear quite insensitive to (i) the *age range* in which income is observed, and (ii) the *number* of income observations available for each person. Such properties would the estimator particularly attractive for comparative purposes, such

[25]The estimator is however still subject to two key limitations. First, we are considering only left-hand side measurement error here, i.e. estimating the lifetime incomes in the child generation. Additional steps would be necessary to adapt the estimator when also the income of parents is not well observed. Second, we assumed that individuals in all age ranges are observed in the estimation of equation (3.12). If the cohorts of interest are not observed over their entire lifecycle, researchers need to approximate the overall shape of their lifecycle profiles via other means. These additional steps are likely to differ across applications, which would then again reduce the comparability of estimates.

Table 3.4: The Revised Lifecycle Estimator

		Direct estimator		Lifecycle estimator			
		Lifetime	Annual	Baseline	Parental	2-step FE (Linear)	2-step FE (Quadratic)
Son's Age	N	(1)	(2)	(3)	(4)	(5)	(6)
Age $\leq$ 27	188,260	0.257***	0.084***	0.165***	0.207***	0.198***	0.213***
		(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.004)
		R2=0.056	R2=0.005	R2=0.023	R2=0.036	R2=0.019	R2=0.017
Age $\leq$ 30	187,432	0.259***	0.129***	0.197***	0.225***	0.232***	0.252***
		(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
		R2=0.058	R2=0.011	R2=0.038	R2=0.049	R2=0.034	R2=0.032
Age $\leq$ 33	186,490	0.259***	0.186***	0.206***	0.223***	0.233***	0.246***
		(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
		R2=0.059	R2=0.022	R2=0.046	R2=0.053	R2=0.042	R2=0.040
Age $\leq$ 36	185,508	0.257***	0.200***	0.215***	0.237***	0.240***	0.251***
		(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
		R2=0.059	R2=0.024	R2=0.050	R2=0.061	R2=0.047	R2=0.045
Age $\leq$ 40	183,886	0.255***	0.235***	0.230***	0.273***	0.261***	0.262***
		(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
		R2=0.059	R2=0.030	R2=0.056	R2=0.078	R2=0.051	R2=0.052
Age $\leq$ 45	181,594	0.254***	0.280***	0.243***	0.274***	0.266***	0.260***
		(0.002)	(0.004)	(0.002)	(0.002)	(0.003)	(0.002)
		R2=0.060	R2=0.031	R2=0.061	R2=0.076	R2=0.058	R2=0.059

Notes: The table reports the slope coefficient from a regression of son's income on father's lifetime income. The measure for son's income is lifetime income in column (1), the annual income at the median age betwee age 25 and the indicated upper age bound in column (2), or the predicted lifetime income from a lifecycle estimator applied to the indicated age range in columns (3) to (6). See main text for a description of each estimator. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

as the comparison of intergenerational estimates across countries. In future extensions we plan to also study its performance in the estimation of mobility trends.

Table 3.5: The Lifecycle Estimator with Few Income Observations

		Lifecycle estimator (2-step FE, Quadratic)				
Son's Age	N	≤ 6 obs.	≤ 5 obs.	≤ 4 obs.	≤ 3 obs.	≤ 2 obs.
Age ≤ 30	187,432	0.242***	0.240***	0.239***	0.240***	0.233***
		(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
		R2=0.033	R2=0.033	R2=0.032	R2=0.031	R2=0.028
Age ≤ 35	185,848	0.244***	0.244***	0.242***	0.237***	0.244***
		(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
		R2=0.041	R2=0.041	R2=0.038	R2=0.035	R2=0.034

Notes: The table reports the slope coefficient from a regression of son's income on father's lifetime income. The measure for son's income is the predicted lifetime income from a lifecycle estimator (2-step FE estimator, interacting the first-stage individual FE with a quadratic in age) applied to the indicated age range. The top row indicates the maximum number of income observations used for each person in the child generation. Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

3.5.2 The Standardized Errors-in-Variables Model

A simple bias correction follows from a *standardized* version of the classical errors-in-variables model. Rewrite annual incomes as

$$y_{st} = \delta_{st}\left(y_s^* + u_{st}\right), \tag{3.13}$$

where δ_{st} is a scaling factor that may vary with age t, while the term in brackets is the classical errors-in-variables model with $Cov(y_s^*, u_{st}) = 0$.[26] Under this errors-in-variables model, the slope coefficient in a regression of (log) annual income for sons $y_{s,t}$ on lifetime income of fathers y_f^*,

$$plim\hat{\beta}_t = \frac{Cov\left(y_{st}, y_f^*\right)}{Var\left(y_f^*\right)} = \beta\delta_{st}, \tag{3.14}$$

is biased by the scaling factor δ_{st}. As noted by Haider and Solon (2006), this factor can be estimated by the slope coefficient in a regression of (log) annual on lifetime income. However, individual-level data containing both annual and lifetime incomes are rarely available, so most

[26] If we assume that the variance of the error u_t does not depend on age this represents a restriction with respect to the *"generalized"* errors-in-variables model proposed by Haider and Solon (2006). This restriction is equivalent to the assumption that the signal-to-noise ratio in annual income as a proxy for lifetime income remains stable over age.

applications simply measure incomes at an age at which the parameter δ_{st} is hoped to be close to one.

We propose to improve on this rule-of-thumb by extracting information from the short snapshots of income that are available in most applications. First, note that the ratio between the variance of annual and lifetime incomes can be expressed as

$$\frac{Var(y_s^*)}{Var(y_{st})} = \frac{Var(y_s^*)}{Var(\delta_{st}\,(y_s^* + u_{st}))} = \frac{1}{\delta_{st}^2}\frac{Var(y_s^*)}{Var(y_s^*) + Var(u_{st})}$$

which in turn implies that

$$\delta_{st} = \left(\frac{Var(y_{st})}{Var(y_s^*)}\right)^{\frac{1}{2}} \left(\frac{Var(y_s^*)}{Var(y_s^*) + Var(u_{st})}\right)^{\frac{1}{2}} \tag{3.15}$$

The bias term δ_{st} can therefore be decomposed into two components, the (i) ratio between the variances of annual and lifetime income, and (ii) the signal-to-noise ratio in annual incomes (or *reliability ratio*). Accordingly, replacing sons' annual incomes $y_{s,t}$ by their standardized values

$$y_{st}^{std} = y_{st}\left(\frac{Var(y_{st})}{Var(y_s^*)}\right)^{-\frac{1}{2}} \left(\frac{Var(y_s^*)}{Var(y_s^*) + Var(u_{st})}\right)^{-\frac{1}{2}}$$

would eliminate the bias in equation (3.14).

This decomposition is useful if these two components can be approximated from short income spans, or be imported from external sources. First consider estimation of the reliability ratio $\frac{Var(y_s^*)}{Var(y_s^*)+Var(u_{st})}$. If the error u_t in equation (3.13) is not serially correlated, this ratio can be estimated by regressing an income observation on its lag (requiring only the observation of two income observations per person). With serial correlation in the transitory component of the income process its estimation becomes more involved, but remains feasible if sufficiently many income observations are available (see e.g. Mazumder 2005). The reliability ratio will therefore be estimable in many applications.[27] Estimation of the ratio between the variances of annual and lifetime incomes is more challenging. However, researchers will often be able to make an educated guess, for example by comparing the variance of annual incomes in the parent and child generation, or can report estimates of the IGE under different assumptions on this variance ratio.

Figure 3.6 illustrates how this standardized errors-in-variables model performs in the

[27] Moreover, the signal-to-noise ratio tends to be relatively stable over age, and can potentially be imported from external sources. For example, Björklund (1993) finds that this ratio (which is equal to the correlation between annual and lifetime incomes) tends to be around two thirds (0.8^2) after age 35 in Swedish sources. Given this stability, external evidence on this ratio might provide a good approximation of its actual size in the data under question. Alternatively, researchers could estimates of the IGE β under a plausible range of the signal-to-noise ratio.

Figure 3.6: A standardized errors-in-variables model in Swedish Data

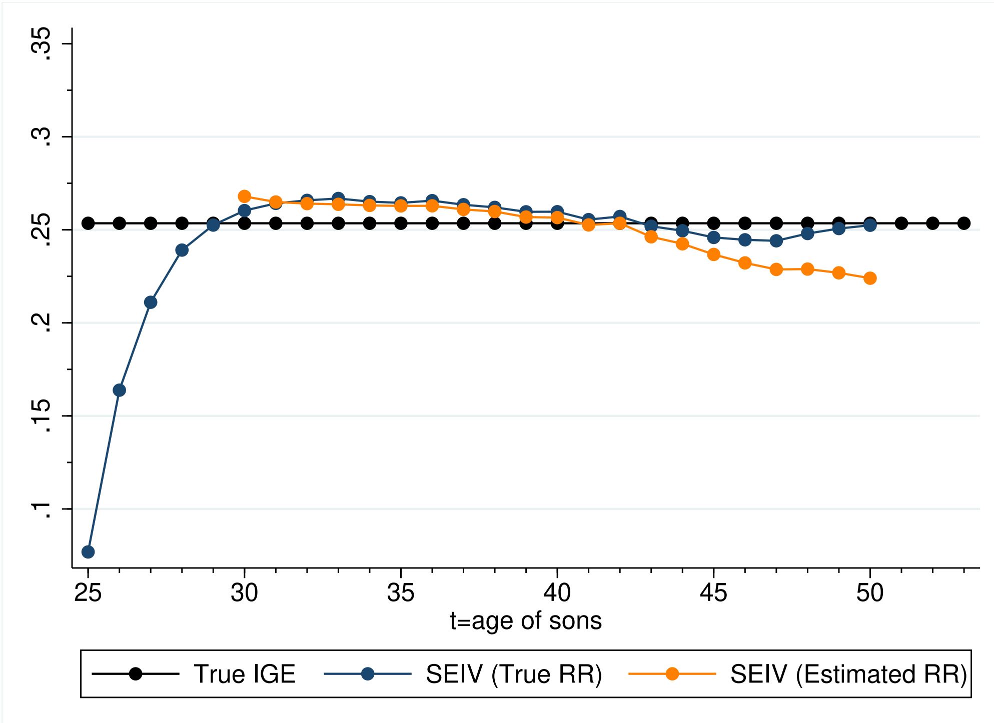

Notes: In this Figure, we implement the standardized errors-in-variables model in the Swedish data. The green line represents the implementation using the true reliability ration. The orange line uses the estimated reliability ration using a regression of one's income on the fourth lagged variable.

Swedish data. We report two bias-corrected series. We abstract from measurement error in $\frac{Var(y_{st})}{Var(y_s^*)}$ by directly estimating it in our data. In the first implementation, we estimate the correlation between annual and lifetime income to derive a "true" estimate of the reliability ratio $\frac{Var(y_s^*)}{Var(y_s^*)+Var(u_{st})}$ (blue line). In a second implementation, we instead estimate the reliability ratio in short income spans, by the correlation between annual income and its fourth lag (orange line). We choose the fourth lag to reduce the influence of serial correlation in transitory shocks. The method performs well over the age range 30 to 50, with either bias-corrected series being within 5-10 percent of the true IGE. Importantly, the bias-correction performs nearly as well when based on a reliability ratio estimated from annual incomes as when based on the true reliability ratio. Only in the late 40s the "true" measure performs measurably better than the "estimated" version.

The standardized errors-in-variables model aims to capture the the strengths of previously proposed correction methods, while addressing its weaknesses. It shares the same motivation as the method proposed by Creedy (1988), that annual income observations should be adjusted by their age-specific variance. However, equation (3.15) clarifies that incomes need to be scaled by both the variance ratio and the reliability ratio, to account for the influence of transitory shocks. Both components are captured by the generalized errors-in-variables model proposed by Haider and Solon (2006), but they can be estimated and corrected for separately.

3.6 Concluding Remarks

The estimation of the intergenerational elasticity (IGE) ideally requires information on the complete income trajectory of parents and offspring. Yet, such ideal datasets are not available for practitioners, who, therefore, use snapshots of income for the IGE estimation. If these snapshots do not mimic lifetime incomes, the resulting estimate would suffer from the so-called lifecycle bias. A large literature has focused on analyzing, discussing and providing some solutions to this measurement issue. One strand of this literature has proposed errors-in-variables models, based on the intuition that there is an age - usually around midlife - in which differences between parental and offspring annual incomes mimic the difference in lifetime income. Most empirical applications (see Table A.1) use this intuition and average income around midlife to reduce lifecycle bias. A second strand of the literature proposes to directly model individuals income trajectories based on incomplete profiles and, then, predict lifetime income.

In this paper, we use nearly complete income profiles of two generations of parents and offspring from Sweden to test how the previously proposed correction methods perform. To do that, we establish a link between the intergenerational and the income process literature, which helps in the assessment of why these methods improve IGE estimates if compared to annual incomes, but still fail to completely eliminate lifecycle bias. Generally, we find that while accounting for important aspects of the income process, each of the methods fail to account for one of its key components: (i) transitory noise, (ii) income growth explained by observable characteristics and (iii) unexplained income growth that nevertheless correlates within families.

Then, we propose two distinct methods to address lifecycle bias in the estimation of the IGE. First, based on the insights from Vogel (2007), Hertz (2007) and Creedy (1988), we introduce a new *lifecycle estimator*. In our preferred version of the estimator, the practitioner should follow two steps. First, estimate an equation of log annual income on age and other available controls (e.g. education, occupation) and on individual fixed effects. Differently from previously proposed methods, we allow for individual slopes to vary systematically by individual. For that, in the second step, we re-estimate the equation incorporating an additional term: an interaction between the predicted individual fixed effect and a quadratic in age. We test our estimator in the Swedish data and find that it works remarkably well.

Estimates are generally within 5 percent of the benchmark IGE, fluctuating around it without any apparent systematic lifecycle bias. The estimator is fairly insensitive to the age range available for the first step estimation and to the number of income observations available for each individual (as long as there is a minimum of two). These properties make the estimator attractive for a large number of applications, and for comparative purposes, both across different countries and across time. In the future, we plan to assess the performance of our estimator in the measurement of mobility trends.

Second, we propose a *standardized errors-in-variables model (SEiV)*, which relates the generalized errors-in-variables model (GEiV) proposed by Haider and Solon (2006) to moments that are more directly obtained from the short income snapshots that are typically observed in practice. Instead of the slope coefficient in a regression of annual on lifetime incomes, as required by the GEiV, the SEiV's adjustment requires two different objects: the ratio between the variance of annual and lifetime income, and the reliability ratio at the age at which incomes are observed. The advantage of this decomposition is that these two components can be more easily approximated from short income spans, or be imported from external sources. Without serial correlation, the reliability ratio can be directly obtained by regressing an income observation on its lag. The ratio between the variances of annual and lifetime income, although more challenging, can, for example, be reported under different assumptions, or under an educated guess. This method also performs well, producing estimates within 5-10 percent of the true IGE.

In sum, both of our proposed methods offer promising alternatives for the measurement problem of the intergenerational literature. Their applicability in different data needs to be studied and tested more thoroughly. However, our results suggest that the methods can be applied in a wide range of settings, and that the IGE can be estimated based on short snapshots of income and simple observable characteristics, such as education or occupation.

Appendix

3.A Additional Figure and Tables

Figure A.1: Components of the Income Process

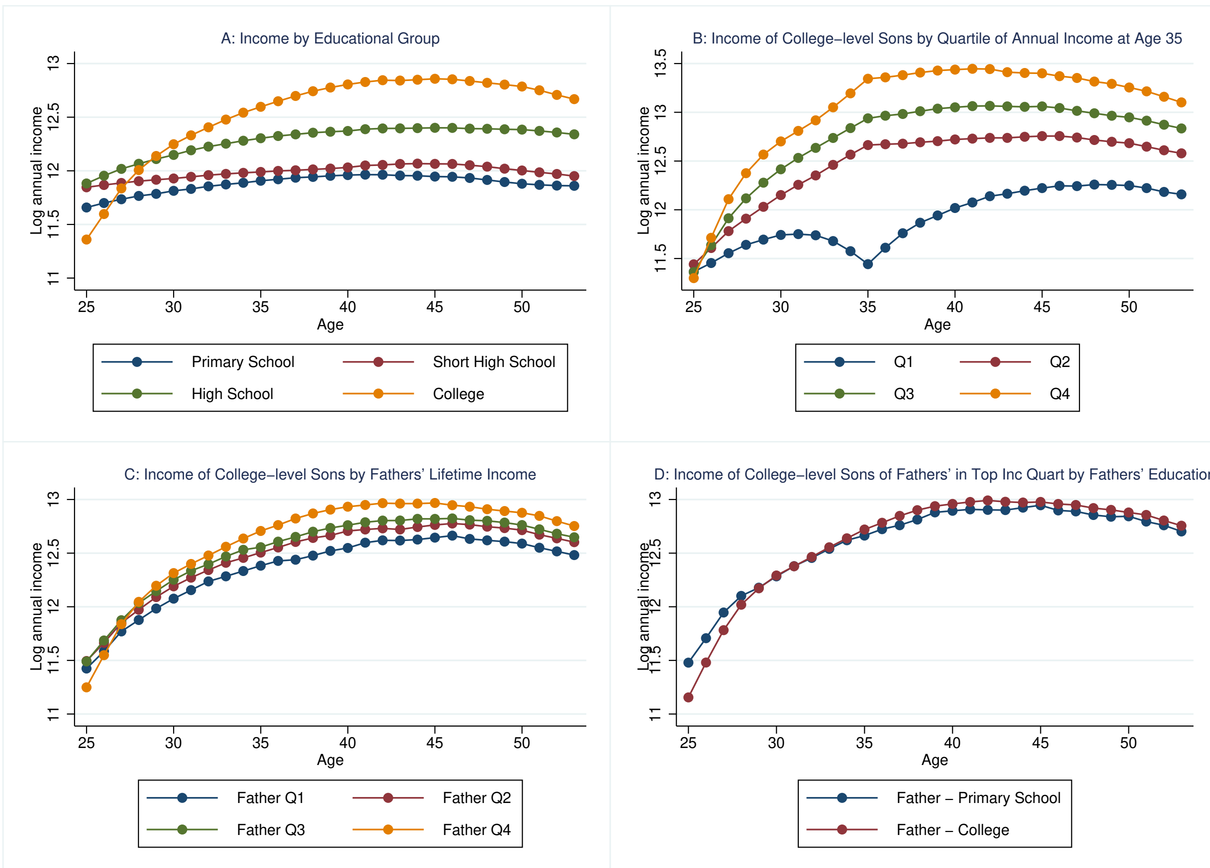

Notes: Panel A shows income trajectories by educational category. Panel B focus only on college-educated sons, who are split in four groups, according to their annual income at age 35. Category Q1 refers to the bottom quartile and Q4 to the top. In Panel C, college-educated sons are divided in four groups, according to fathers' lifetime income. Finally, in Panel D, college-level sons whose fathers belong to the top quartile of lifetime income are divided in two additional groups: college-educated fathers and fathers with only primary school.

Table A.1: Intergenerational Elasticity Literature

Authors	Year	Journal	Estimate (US)	Method	Details on approach to attenuate lifecycle bias on offspring
IGE in Levels					
Solon	1992	AER	0.41	Averaging	Single year annual earnings, average age 29.6
Zimmerman	1992	AER	0.54	Averaging	Single year of son's annual earnings; average age 33.8
Mazumder	2005a	Restat	0.61	Averaging	Average 4 years of income, age 30-35
Mazumder	2005b	Book chapter	0.62	Averaging	Single year of son's income, but implements HEIV
Hertz	2006	Working Paper	0.58	Averaging	Average 4.1 income observations at mean-age 37
Bratsberg et al	2007	Economic Journal	0.54	Averaging	Annual income in 1995 and 2001, cohorts 1957-1964.
Gouskova et al	2010	Labour Economics	0.63	Averaging	Single year from ages 35-44
Chau et al	2012	Economic Letters	0.66	Model Income	At least 3 observations of annual earnings between the ages of 25-60. Use earnings dynamics model.
Janti and Lindahl	2012	Economic Letters	Sweden	Model Income	Formulate simple model with heterogeneous income
Chetty et al	2014	QJE	0.34	Averaging	2-year average around ages 29-32 (2011 and 2012)
Mitnik et al	2015	Working Paper	0.56	Averaging	Single year around ages 35-38
Mazumder	2016	Research in Labor Economics	0.66	Averaging	Average between 1 and 11 years around age 40
Borisov & Pissarides	2016	IZA WP	Russia	Model Income	Predicted value of permanent earnigns based on monthly earnings. Controls for hours worked, age, year of birth, education.
Landers & Heckman	2017	Scandinavian Journal	0.29 to 0.45	Averaging	Average between ages 34-41 until 30-35
IGE in Trends					
Mayer & Lopoo	2005	Journal of Human Resources	Non-linear	Averaging	Son's family income at age 30
Hertz	2007	Industrial Relations	No trend	Model Income	Hertz correction
Aaronson & Mazumder	2008	Journal of Human Resources	Non-linear	-	
Lee & Solon	2009	Restat	No trend	Averaging	Average of all available years, changing across cohorts
Justman & Krush	2013	Working Paper	Upward	Model Income	Predicted income at age 40, controls for age, education, race, marital status and individual FE
Justman et al	2017	Working Paper			
Hartley et al	2017	Working Paper	Upward	Averaging	Multi-year average, Lee & Solon age adjustment (mother-daughters)
Davis & Mazumder	2017	Working Paper	Upward	Averaging	Average of 3 year of son's family income
Lifecycle Bias (Methodological Papers)					
Haider & Solon	2006	AER	-	-	Proposal GEiV
Grawe	2006	Labour Economics	-	-	Discussion of lifecycle bias
B́ohlmark & Lindquist	2006	Journal of Labor Economics	-	-	Discussion of lifecycle and GEiV
Nilsen et al	2012	Scandinavian Journal	-	-	Discussion of lifecycle bias
Nybom & Stuhler	2016	Journal of Human Resources	-	-	Testing lifecycle bias and GEiV
Chen et al.	2017	Labour Economics	-	-	Discussion of life-cycle

Table A.2: Moments of Income Process in Guvenen (2009)

Moment	Value
ρ	0.821
σ^2_α	0.022
σ^2_β	0.00038
$corr_{\alpha\beta}$	-0.23
σ^2_η	0.029
σ^2_ε	0.047

Figure A.2: Variance of Income by Group

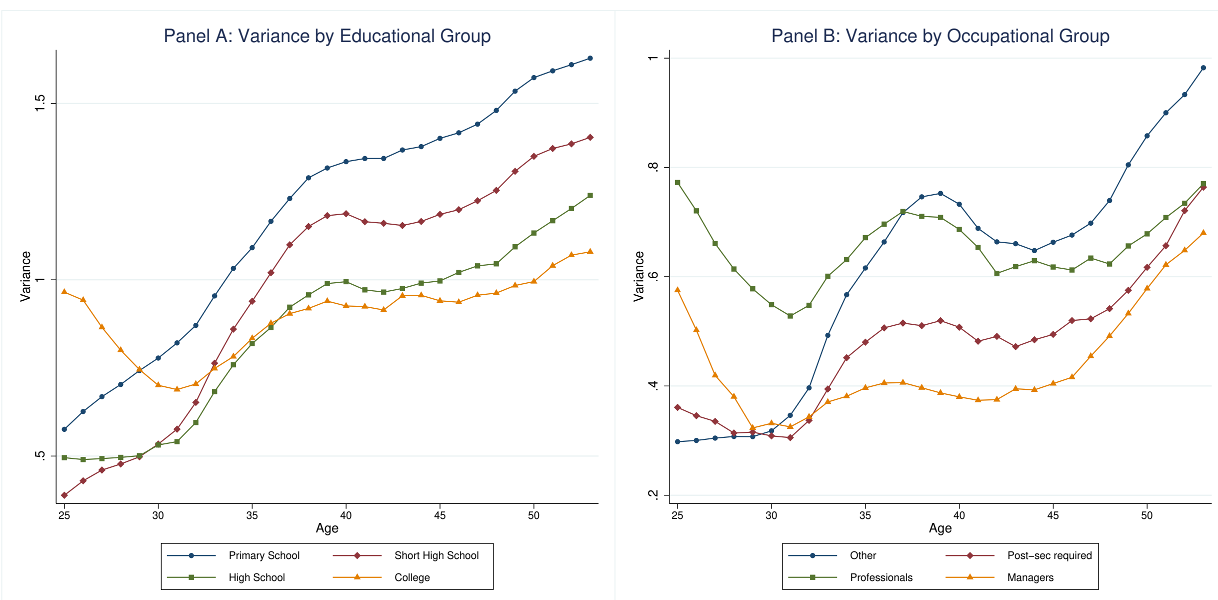

Notes: In this Figure, we plot how the variance of income varies by age, in the Swedish data. In Panel A, we graph it for different educational groups and, in Panel B, for different occupational groups.

Bibliography

An, Y., L. Wang, and R. Xiao (2017). Your American Dream is Not Mine! A New Approach to Estimating Intergenerational Mobility Elasticities. *Working Paper*.

Antonovics, K. and B. Backes (2014). The effect of banning affirmative action on human capital accumulation prior to college entry. *IZA Journal of Labor Economics 3*(1), 1–20.

Arcidiacono, P. and M. Lovenheim (2016). Affirmative Action and the Quality–Fit Trade-off. *Journal of Economic Literature 54*(1), 3–51.

Arcidiacono, P., M. Lovenheim, and M. Zhu (2015). Affirmative Action in Undergraduate Education. *Annual Review of Economics 7*(1), 487–518.

Assouad, L., L. Chancel, and M. Morgan (2018). Extreme Inequality: Evidence from Brazil, India, the Middle East, and South Africa. *AEA Papers and Proceedings 108*(January), 119–23.

Assunção, J. and B. Ferman (2015). Does affirmative action enhance or undercut investment incentives? Evidence from quotas in Brazilian public universities. *Working Paper*.

Atkinson, A. (1980). On Intergenerational Income Mobility in Britain. *Journal of Post Keynesian Economics 3*(2), 194–218.

Autor, D. (2003). Outsourcing at Will. *Journal of Labor Economics 21*(1), 1–42.

Backes, B. (2012). Do affirmative action bans lower minority college enrollment and attainment? Evidence from statewide bans. *Journal of Human Resources 47*(2), 435–455.

Bagde, S., D. Epple, and L. Taylor (2016). Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India. *American Economic Review 106*(6), 1495–1521.

Bertrand, M., R. Hanna, and S. Mullainathan (2010). Affirmative action in education: Evidence from engineering college admissions in India. *Journal of Public Economics 94*(1-2), 16–29.

Bettinger, E. P., B. Terry Long, P. Oreopoulos, and L. Sanbonmatsu (2012). Information in College Decisions: Results From The H&R Block FAFSA Experiment. *The Quarterly Journal of Economics* (January), 1205–1242.

Björklund, A. (1993). A comparison between actual distributions of annual and lifetime income: Sweden 1951-89. *Review of Income and Wealth 39*(4).

Black, S. and P. Devereux (2011). Recent Developments in Intergenerational Mobility. *Handbook of Labor Economics 4*(B), 1487–1541.

Blanden, J. (2011). Cross-country rankings in intergenerational mobility: a comparison of approaches from economics and sociology. *Journal of Economic Surveys 27*(1), 38–73.

Bó, I. and R. Hakimov (2016). The Iterative Deferred Acceptance Mechanism. *SSRN 49*(November), 1–61.

Bodoh-Creed, A. and B. Hickman (2018). Pre College Human Capital Investment and Affirmative Action: a Structural Policy Analysis of US College Admissions. *Working Paper*.

Böhlmark, A. and M. J. Lindquist (2006). Life-Cycle Variations in the Association between Current and Lifetime Income: Replication and Extension for Sweden. *Journal of Labor Economics 24*(4), 879–896.

Bulman, G. (2015). The effect of access to college assessments on enrollment and attainment. *American Economic Journal: Applied Economics 7*(4), 1–36.

Carneiro, P., I. Garcia, K. Salvanes, and E. Tominey (2015). Intergenerational Mobility and the Timing of Parental Income. *Ssrn* (October).

Carrell, S. and B. Sacerdote (2017). Why Do College-Going Interventions Work? *American Economic Journal: Applied Economics 9*(3), 124–151.

Chau, T. W. (2012). Intergenerational income mobility revisited: Estimation with an income dynamic model with heterogeneous age profile. *Economics Letters 117*(3), 770–773.

Chetty, R., J. Friedman, E. Saez, N. Turner, and D. Yagan (2017). Mobility Report Cards: The Role of Colleges in Intergenerational Mobility.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quaterly Journal of Economics 129*(4), 1553–1623.

Chetty, R., N. Hendren, P. Kline, E. Saez, and N. Turner (2014). Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility. *American Economic Review P&P 104*(5), 141–147.

Corak, M. (2013). Income Inequality, Equality of Opportunity, and Intergenerational Mobility. *Journal of Economic Perspectives 27*(3), 79–102.

Cotton, C. S., B. R. Hickman, and J. P. Price (2018). Incentive Provision in Investment Contests: Theory and Evidence. *Working Paper*.

Creedy, J. (1988). Earnings Comparisons Between Generations: Some Alternative Approaches. *The Manchester School 56*(3), 268–281.

Cullen, J. B., M. C. Long, and R. Reback (2013). Jockeying for position: Strategic high school choice under Texas' top ten percent plan. *Journal of Public Economics 97*(1), 32–48.

Davis, J. and B. Mazumder (2019). The Decline in Intergenerational Mobility After 1980. *Federal Reserve Bank of Chicago Working Paper*.

Deming, D. and S. Dynarski (2010). Into College, Out of Poverty? Policies to Increase the Postsecondary Attainment of the Poor. In *Targeting Investments in Children: Fighting Poverty When Resources Are Limited*, Number September. University Chicago Press.

Denning, J. T. (2017). College on the Cheap: Costs and Benefits of Community College. *American Economic Journal: Economic Policy 9*(2), 155–188.

Dillon, E. W. and J. A. Smith (2017). Determinants of the Match between Student Ability and College Quality. *Journal of Labor Economics 35*(1), 45–66.

Estevan, F., T. Gall, P. Legros, and A. Newman (2019). The Top-Ten Way to Integrate High Schools. *Working Paper*.

Estevan, F., T. Gall, and L.-P. Morin (2019). Redistribution without distortion: Evidence from an affirmative action program at a large Brazilian university. *Economic Journal (Forthcoming)*.

Francis, A. and M. Tannuri-Pianto (2018). Black Movement: Using discontinuities in admissions to study the effects of college quality and affirmative action. *Journal of Development Economics 135*(September 2017), 97–116.

Francis, A. M. and M. Tannuri-Pianto (2012). Using Brazil's Racial Continuum to Examine the Short-Term Effects of Affirmative Action in Higher Education. *Journal of Human Resources 47*(3), 754–784.

Goldhaber, D. and G. Peri (2007). Community Colleges. In *Economic Inequality and Higher Education: Access, Persistence and Success*, pp. 101–127. Russel Stage Foundation.

Gottschalk, P. and R. Moffitt (1994). The Growth of Earnings Instability in the U.S. Labor Market. *Brookings Papers on Economic Activity 25*(2), 217–272.

Grawe, N. D. (2006). Lifecycle bias in estimates of intergenerational earnings persistence. *Labour Economics 13*, 551–570.

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

Haider, S. (2001). Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991. *Journal of Labor Economics 19*(4), 799–836.

Haider, S. and G. Solon (2006). Variation in the Association between Current and Life-Cycle Lifetime Earnings. *American Economic Review 96*(4), 1308–1320.

Haveman, R. and K. Wilson (2007). Access, Matriculation, and Graduation. In *Economic Inequality and Higher Education: Access, Persistence and Success*, pp. 17–43. Russel Stage Foundation.

Heidrich, S. (2016). A study of the missing data problem for intergenerational mobility using simulations s. *Working Paper*.

Hertz, T. (2007). Trends in the intergenerational elasticity of family income in the United States. *Industrial Relations 46*(1), 22–50.

Hinrichs, P. (2012). The Effects of Affirmative Action Bans on College Enrollment, Educational Attainment, and the Demographic Composition of Universities. *Review of Economics and Statistics 94*(3), 712–722.

Hoxby, C. M. and S. E. Turner (2015). What High-Achieving Low-Income Students Know About College. *American Economic Review: Papers & Proceedings 105*(5), 514–517.

Hryshko, D. (2012). Labor income profiles are not heterogeneous: Evidence from income growth rates. *Quantitative Economics 3*(2), 177–209.

Jäntti, M. and S. Jenkins (2015). *Income Mobility* (1 ed.), Volume 2. Elsevier B.V.

Jäntti, M. and L. Lindahl (2012). On the variability of income within and across generations. *Economics Letters 117*(1), 165–167.

Jenkins, S. (1987). Snapshots versus movies: 'Lifecycle biases' and the estimation of intergenerational earnings inheritance. *European Economic Review 31*(5), 1149–1158.

Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quaterly Journal of Economics* (May).

Kaufmann, K. M. (2014). Understanding the income gradient in college attendance in Mexico: The role of heterogeneity in expected returns. *Quantitative Economics 5*(3), 583–630.

Kitagawa, T., M. Nybom, and J. Stuhler (2019). Measurement error and rank correlations. *Working Paper*.

Lee, C. and G. Solon (2009). Trends in Intergenerational Income Mobility. *Review of Economics and Statistics 91*(4), 766–772.

Machado, C. and C. Szerman (2018). Centralized Admission and the Student-College Match. *Working Paper*.

MaCurdy, T. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics 18*(1), 83–114.

Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the united states using social security earnings data. *Review of Economics and Statistics 87*(2), 235–255.

Mazumder, B. (2016). Estimating the intergenerational elasticity and rank association in the United States: Overcoming the current limitations of Tax data. *Research in Labor Economics 43*(September), 83–129.

Meghir, C. and L. Pistaferri (2011). Earnings, Consumption and Lifecycle Choices. *Handbook of Labor Economics 4*(Part B Chapter 9).

Moffitt, R. and P. Gottschalk (1995). Trends in the Covariance Structure of Earnings in the U.S. 1969 –1987. *Institute for Research on Poverty Discussion Paper 1001-93*.

Nilsen, O., K. Vaage, A. Aakvik, and K. Jacobsen (2012). Intergenerational Earnings Mobility Revisited: Estimates Based on Lifetime Earnings. *Scandinavian Journal of Economics 114*(1), 1–23.

Nybom, M. and J. Stuhler (2016). Heterogeneous Income Profiles and Lifecycle Bias in Intergenerational Mobility Estimation. *Journal of Human Resources 51*(1), 239–268.

Nybom, M. and J. Stuhler (2017). Biases in Standard Measures of Intergenerational Income Dependence. *Journal of Human Resources 52*(3), 800–825.

OECD (2017). Education at a Glance 2017: OECD Indicators. Technical report, OECD Publishing, Paris.

Oreopoulos, P. and R. Dunn (2013). Information and College Access: Evidence from a Randomized Field Experiment. *Scandinavian Journal of Economics 115*(1), 3–26.

Pallais, A. (2015). Small Differences That Matter: Mistakes in Applying to College. *Journal of Labor Economics 33*(2), 493–520.

Rodriguez, J., S. Urzua, and L. Reyes (2016). Heterogeneous Economic Returns to Post-Secondary Degrees: Evidence from Chile. *Journal of Human Resources 51*(2), 416–460.

Rojas, E., T. Rau, and S. Urzúa (2013). Loans for Higher Education: Does the Dream Come True? *NBER Working Paper 19138*.

Solis, A. (2017). Credit Access and College Enrollment. *Journal of Political Economy 125*(2), 562–622.

Solon, G. (1999). Intergenerational Mobility in the Labor Market. *Handbook of Labor Economics 3*(A), Chapter 29.

Vogel, T. (2007). Reassessing intergenerational mobility in Germany and the United States: the impact of differences in lifecycle earnings patterns. *Working Paper* (June).

Wooldridge, J. (2006). *Introductory econometrics: a modern approach* (6th editio ed.). Mason, OH: Thomson/South-Western.

Zimmerman, S. D. (2014). The Returns to College Admission for Academically Marginal Students. *Journal of Labor Economics 32*(4), 711–754.